Institute Henri Poincare
January 2014

Western University
Canada

Higher-order Segmentation Functionals:

Entropy, Color Consistency, Curvature, etc.

Yuri Boykov

jointly with


M. Tang


I. Ben Ayed


O. Veksler

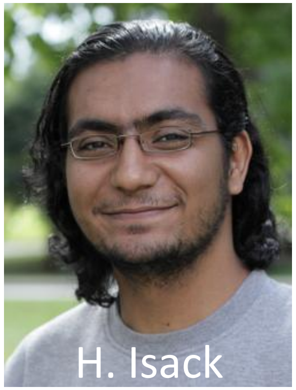
H. Isack


L. Gorelick


C. Nieuwenhuis


E. Toppe


C. Olsson


A. Osokin


A. Delong

Different surface representations

continuous optimization
$s_p \in \Re$

combinatorial optimization
$s_p \in \{0,1\}$

mixed optimization
$s_p \in \mathbf{Z} \quad \theta_p \in \Re$

mesh

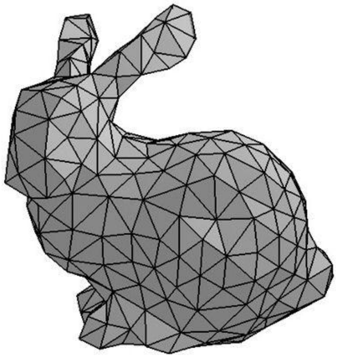

level-sets

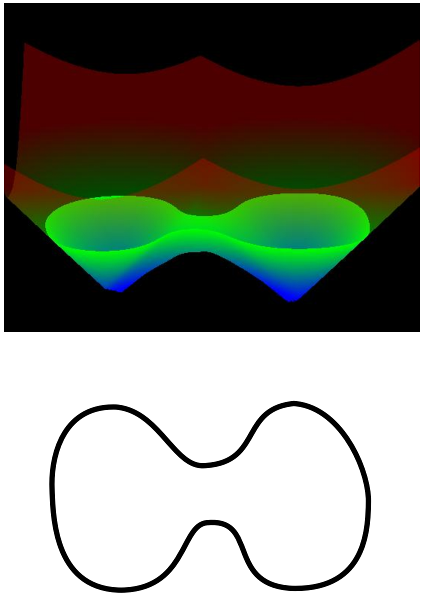

graph labeling

on complex

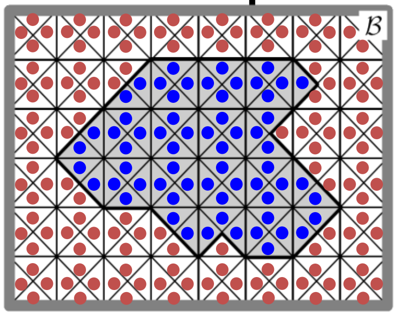

on grid

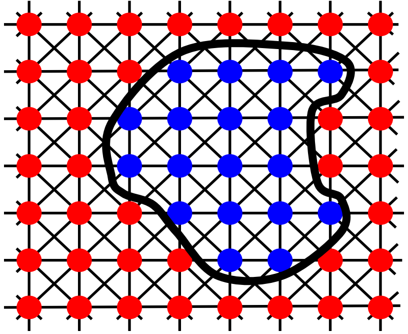

point cloud labeling

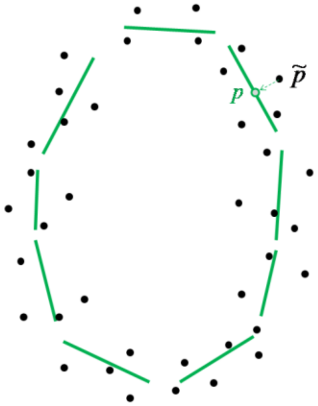

this talk

combinatorial
optimization

$$s_p \in \{0,1\}$$

graph labeling

on grid

Implicit surfaces/bondary ⟶

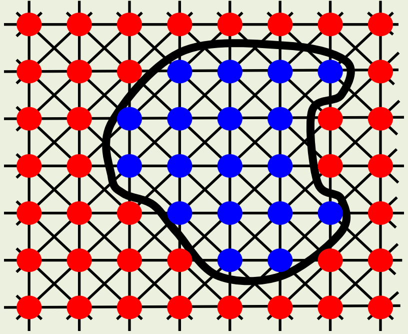

Image segmentation

Basics

$$s_p \in \{0,1\}$$

$$E(S) = \langle f, S \rangle + B(S)$$

$$E(S) = \sum_p f_p s_p + B(S)$$

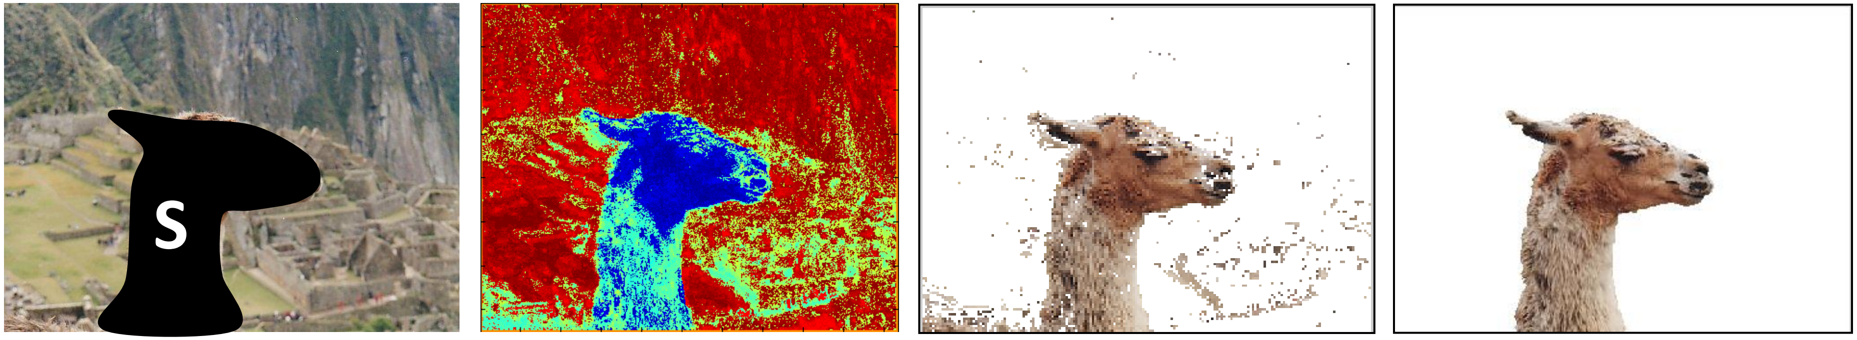

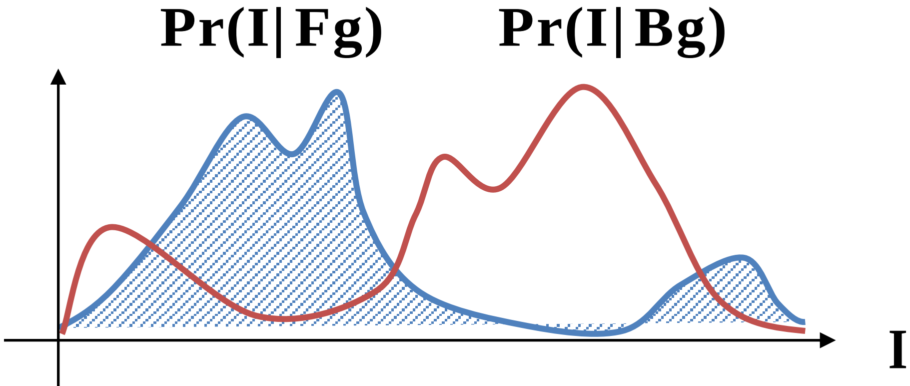

$$f_p = -\ln\left(\frac{\Pr(I_p \mid fg)}{\Pr(I_p \mid bg)}\right)$$

Linear (modular) appearance of **region** S

$$R(S) = \langle f, S \rangle = \sum_p f_p \cdot s_p$$

Examples of potential functions $\boldsymbol{f}$

- Log-likelihoods $f_p = -\ln \Pr(I_p)$
- Chan-Vese $f_p = (I_p - c)^2$
- Ballooning $f_p = -\boldsymbol{1}$

Basic **boundary** regularization for S

$$B(S) = \sum_{pq \in N} w \quad \cdot [s_p \neq s_q]$$

$s_p \in \{0,1\}$

pair-wise discontinuities

Basic **boundary** regularization for S

$$B(S) = \sum_{pq \in N} w \quad \cdot [s_p \neq s_q]$$

$$s_p \in \{0,1\}$$

second-order terms

$$[s_p - s_q] \quad = \quad s_p \cdot (1 - s_q) \; + \; (1 - s_p) \cdot s_q$$

quadratic

Basic **boundary** regularization for S

$$B(S) = \sum_{pq \in N} w_{pq} \cdot [s_p \neq s_q]$$

$s_p \in \{0,1\}$

second-order terms

Examples of discontinuity penalties $\boldsymbol{w}$

- Boundary length $w_{pq} = 1$
- Image-weighted boundary length $w_{pq} = exp(I_p - I_q)^2$

Basic **boundary** regularization for S

$$B(S) = \sum_{pq \in N} w_{pq} \cdot [s_p \neq s_q]$$

$s_p \in \{0,1\}$

second-order terms

- corresponds to boundary length $|\partial S|$
  - grids [B&K, 2003], via integral geometry
  - complexes [Sullivan 1994]
- submodular **second-order** energy
  - can be minimized exactly via *graph cuts*

    [Greig et al.'91, Sullivan'94, Boykov-Jolly'01]

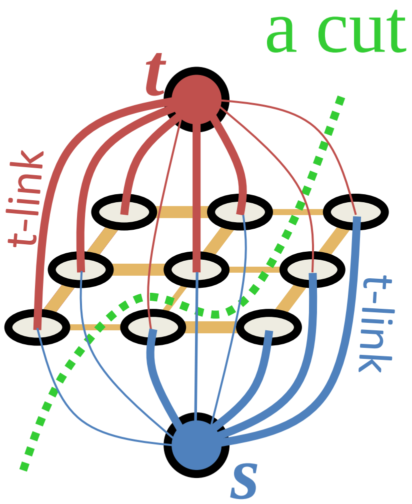

Submodular set functions

any (binary) segmentation energy *E(S)* is a set function *E:* $2^{\Omega} \rightarrow \Re$

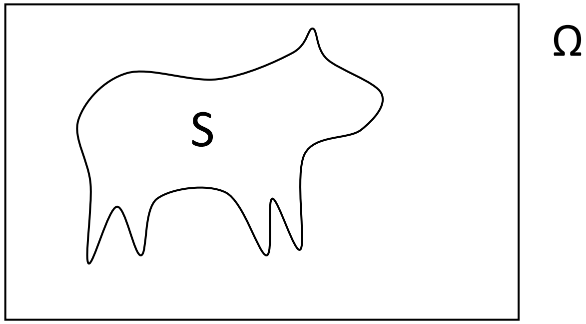

Submodular set functions

Set function $\mathrm{E}:2^{\Omega}\rightarrow\Re$ is **submodular** if for any $S,T\subseteq\Omega$

$$E(S\bigcap T)+E(S\bigcup T)\leq E(S)+E(T)$$

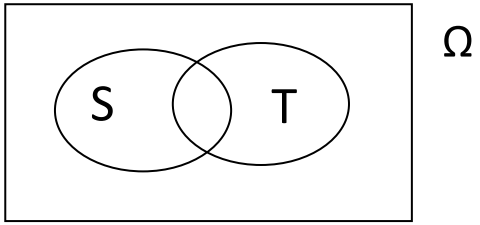

Significance: any submodular set function can be globally optimized in polynomial time [Grotschel et al.1981,88, Schrijver 2000]

$O(|\Omega|^{9})$

Submodular set functions

an alternative equivalent definition providing intuitive interpretation: “**diminishing returns**”

Set function $\mathrm{E}: 2^{\Omega} \to \Re$ is **submodular** if for any $S \subseteq T \subseteq \Omega$

$$E(T \cup \{v\}) - E(T) \le E(S \cup \{v\}) - E(S)$$

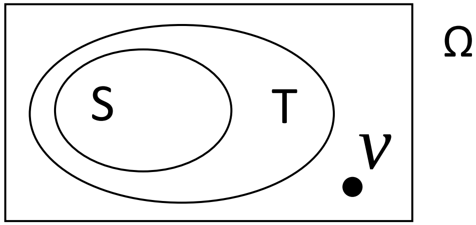

$$\forall v \in \Omega$$

Easily follows from the previous definition: $E(\overbrace{T \cup \{v\}}^{S' \cup T}) + E(\overbrace{S}^{S' \cap T}) \le E(\overbrace{S \cup \{v\}}^{S'}) + \mathrm{E(T)}$

Significance: any submodular set function can be globally optimized in polynomial time [Grotschel et al.1981,88, Schrijver 2000]

$$O(|\Omega|^9)$$

Graph cuts for minimization of submodular set functions

Assume set Ω and *2nd*-order (quadratic) function

$$E(s) = \sum_{(pq)\in N} E_{pq}(s_p, s_q)$$

$$s_p, s_q \in \{\boldsymbol{0}, \boldsymbol{1}\}$$

Indicator variables

Function *E(S)* is **submodular** if for any $(p,q)\in N$

$$E_{pq}(\mathbf{0},\mathbf{0}) + E_{pq}(\mathbf{1},\mathbf{1}) \le E_{pq}(\mathbf{1},\mathbf{0}) + E_{pq}(\mathbf{0},\mathbf{1})$$

Significance: submodular 2nd-order boolean (set) function can be globally optimized in polynomial time by **graph cuts**

[Hammer 1968, Pickard&Ratliff 1973] $O(|\mathrm{N}|\cdot|\Omega|^2)$

[Boros&Hammer 2000, Kolmogorov&Zabih2003]

Global Optimization

Combinatorial optimization

Continuous optimization

submodularity

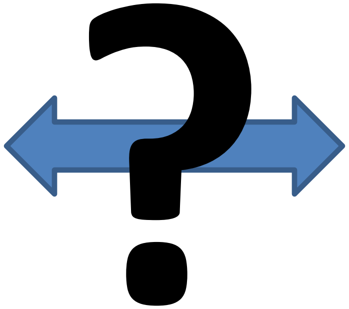

convexity

Graph cuts for minimization of posterior energy (MRF)

Assume **Gibbs distribution** over binary random variables $s_p \in \{\boldsymbol{0},\boldsymbol{1}\}$

$$Pr(s_1,\dots,s_n) \quad \propto \; exp(-E(S)) \qquad \text{for} \qquad S = \{\ p \mid s_p = 1\ \}$$

Theorem [*Boykov, Delong, Kolmogorov, Veksler* in unpublished book 2014?]

All random variables s_p are **positively correlated** iff set function $E(S)$ is **submodular**

That is, submodularity implies MRF with "smoothness" prior

Basic segmentation energy

$$\sum_{p} f_p \cdot s_p \quad + \quad \sum_{pq \in N} w_{pq} \cdot [s_p \neq s_q]$$

segment region/appearance **boundary smoothness**

Higher-order binary segmentation

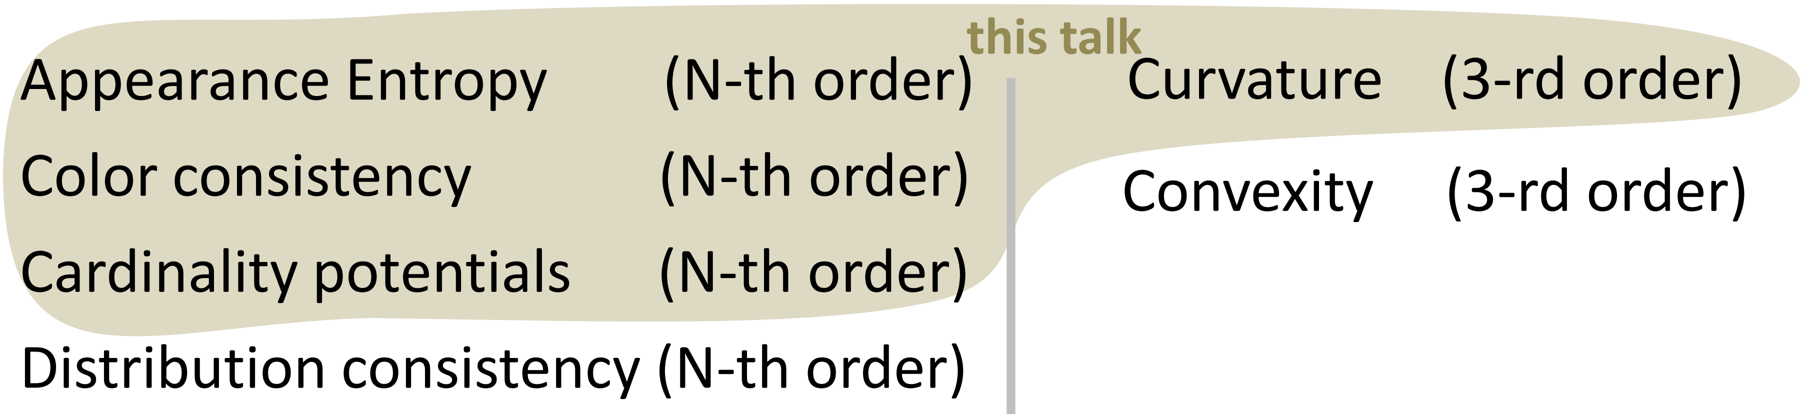

segment region/appearance

Appearance Entropy (N-th order)
Color consistency (N-th order)
Cardinality potentials (N-th order)
Distribution consistency (N-th order)

boundary smoothness

Curvature (3-rd order)
Convexity (3-rd order)

Connectivity (N-th order)
Shape priors (N-th order)

Overview of this talk

high-order functionals	optimization
• From likelihoods to **entropy**	**block-coordinate descent** [Zhu&Yuille 96, GrabCut 04]
• From entropy to **color consistency**	**global minimum** [our work: One Cut 2014]
• Convex **cardinality potentials** • **Distribution consistency** • From length to **curvature**	**submodular approximations** [our work: Trust Region 13, Auxiliary Cuts 13]

other extensions [arXiv13]

Given likelihood models

unary (linear) term

pair-wise (quadratic) term

$$E(S\,|\,\theta_0,\theta_1) \;=\; \sum_{p} -\ln \Pr(I_p\,|\,\theta_{s_p}) \;+\; \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q] \qquad s_p\in\{0,1\}$$

assuming known

- parametric models – e.g. Gaussian or GMM
- non-parametric models - histograms

guaranteed globally optimal S

$I_p \in RGB$

image segmentation, graph cut

[Boykov&Jolly, ICCV2001]

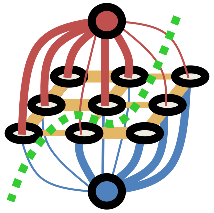

Beyond fixed likelihood models

mixed optimization term   pair-wise (quadratic) term

$$E(S,\theta_0,\theta_1) \;=\; \sum_p -\ln \Pr(I_p \mid \theta_{s_p}) \;+\; \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q] \qquad s_p \in \{0,1\}$$

extra variables

- parametric models – e.g. Gaussian or GMM
- non-parametric models - histograms

NP hard mixed optimization!

[Vesente et al., ICCV'09]

$I_p \in RGB$

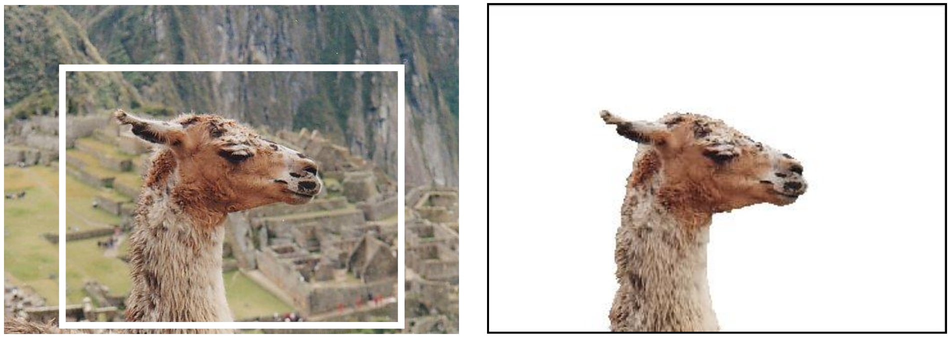

Models θ_0 , θ_1 are iteratively re-estimated (from initial box)

iterative image segmentation, Grabcut

(block coordinate descent $S \leftrightarrow \theta_0,\theta_1$)

[Rother, et al. SIGGRAPH'2004]

Block-coordinate descent for $E(S,\theta_0,\theta_1)$

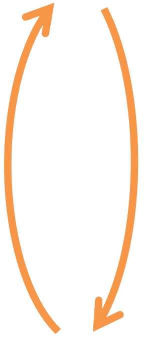

- Minimize over segmentation S for fixed θ_0 , θ_1

$$E(S,\cancel{\theta_0,\theta_1}) = \sum_p -\ln\Pr(I_p|\theta_{s_p}) + \sum_{pq\in N} w_{pq}\cdot[s_p\neq s_q]$$

optimal S is computed using graph cuts, as in [BJ 2001]

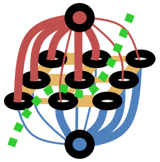

- Minimize over θ_0 , θ_1 for fixed labeling S

$$E(\cancel{S},\theta_0,\theta_1) = \sum_{p:s_p=0} -\ln\Pr(I_p|\theta_0) + \sum_{p:s_p=1} -\ln\Pr(I_p|\theta_1) + \overset{\text{fixed for } S=const}{\sum_{pq\in N} w_{pq}\cdot[s_p\neq s_q]}$$

$$\Downarrow \qquad\qquad \Downarrow$$

$$\hat{\theta}_0 = p^{\bar{S}}$$

distribution of intensities in current bkg. segment $\bar{S}$ $=\{p:S_p=0\}$

$$\hat{\theta}_1 = p^{S}$$

distribution of intensities in current obj. segment $S=\{p:S_p=1\}$

Iterative learning of color models

(binary case $s_p \in \{0,1\}$)

- GrabCut: <u>iterated</u> graph cuts [Rother et al., SIGGRAPH 04]

$$E(S,\theta_0,\theta_1) = \sum_p -\ln \Pr(I_p | \theta_{S_p}) + \sum_{pq \in N} w_{pq} \cdot [s_p \neq s_q]$$

start from models θ_0 , θ_1
inside and outside some given box

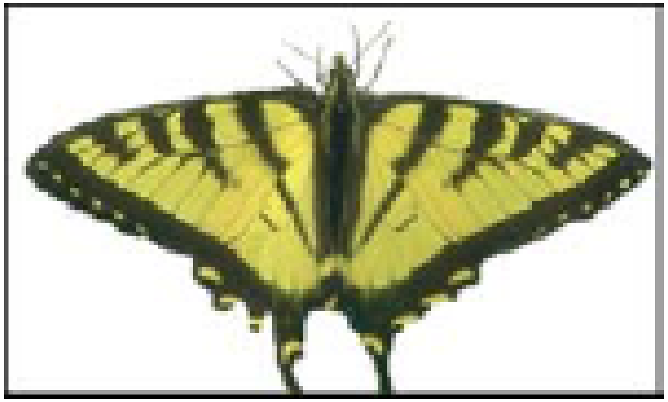

iterate graph cuts and model re-estimation
until convergence to **a local minimum**

solution is sensitive to initial box

Iterative learning of color models

(binary case $s_p \in \{0,1\}$)

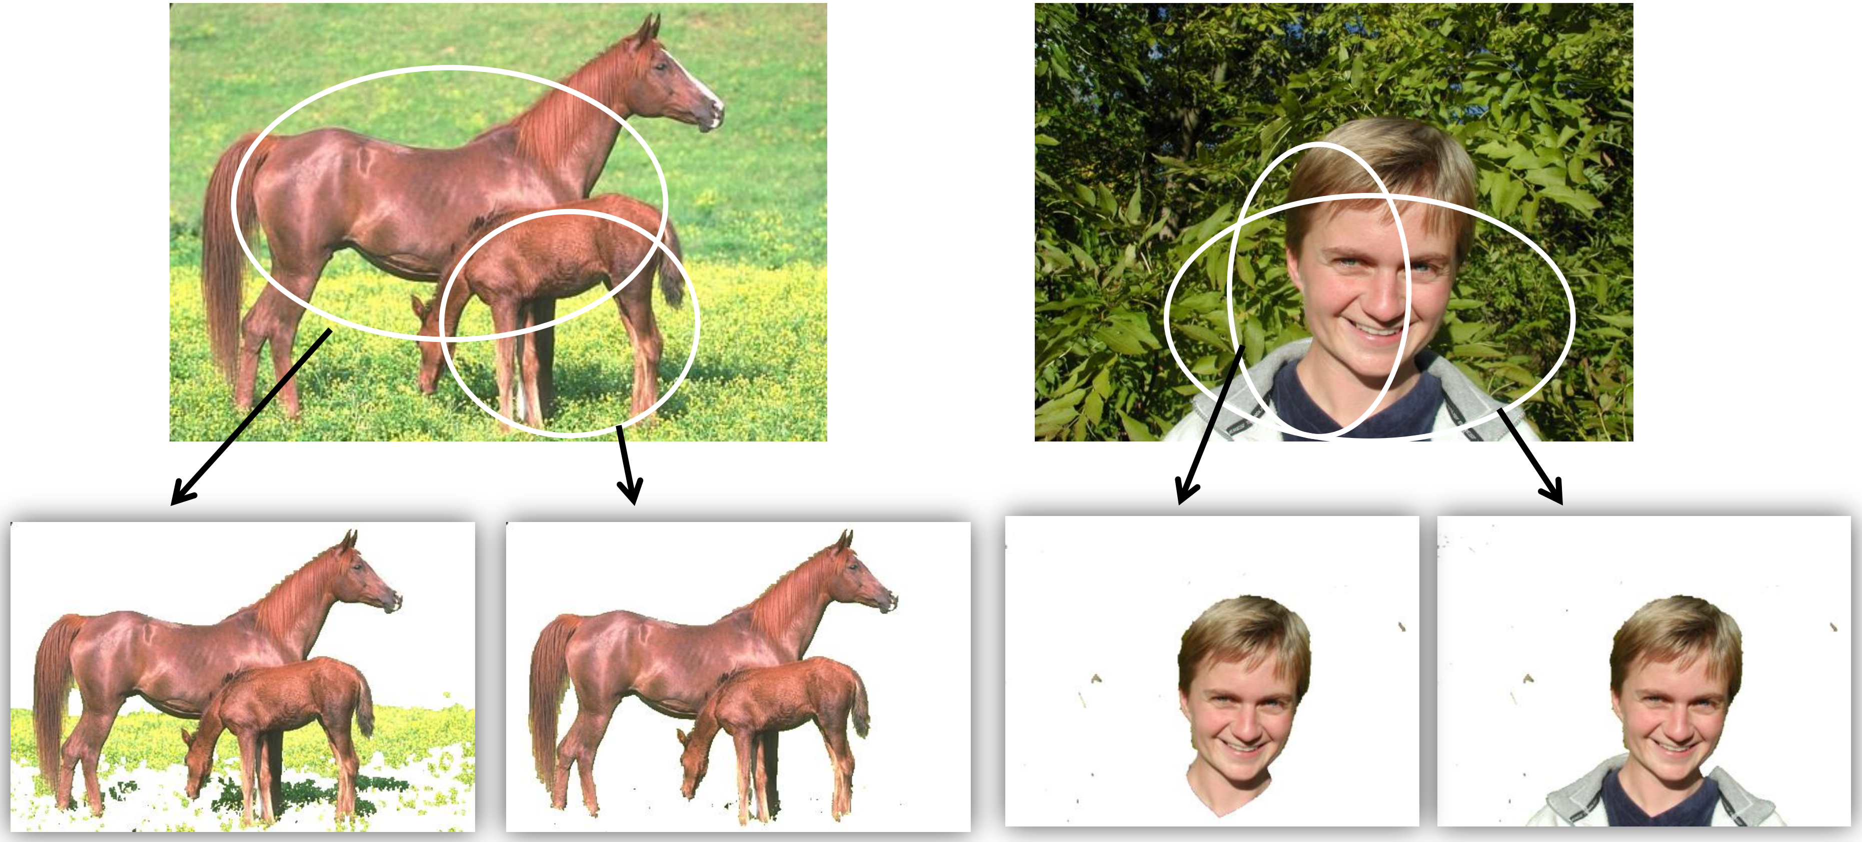

$E=1.410\times10^6$ $E=1.39\times10^6$ $E=2.41\times10^6$ $E=2.37\times10^6$

BCD minimization of $E(S,\theta_0,\theta_1)$ converges to a local minimum

(interactivity a la “snakes”)

Iterative learning of color models

(could be used for more than 2 labels $s_p \in \{0,1,2,...\}$)

- Unsupervised segmentation [Zhu&Yuille, 1996]

using level sets + merging heuristic

$$E(S,\theta_0,\theta_1,\theta_2...) = \sum_p -\ln \Pr(I_p \mid \theta_{S_p}) + \sum_{pq\in N} w_{pq} \cdot [s_p \neq s_q] + |labels|$$

initialize models θ_0 , θ_1 , θ_2 , ...
from many randomly sampled boxes

iterate segmentation
and model re-estimation
until convergence

models compete, stable result if sufficiently many

Iterative learning of color models

(could be used for more than 2 labels $s_p \in \{0,1,2,...\}$)

- Unsupervised segmentation [Delong et al., 2012]

using a-expansion (graph-cuts)

$$E(S,\theta_0,\theta_1,\theta_2 ...) \;=\; \sum_p -\ln \Pr(I_p \,|\, \theta_{S_p}) \;+\; \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q] \;+\; |labels|$$

initialize models θ_0 , θ_1 , θ_2 , …
from many randomly sampled boxes

iterate segmentation
and model re-estimation
until convergence

models compete, stable result if sufficiently many

Iterative learning of other models

(could be used for more than 2 labels $s_p \in \{0,1,2,...\}$)

- Geometric multi-model fitting [Isack et al., 2012]

using a-expansion (graph-cuts)

$$E(S,\theta_0,\theta_1,\theta_2 ...) = \sum_p \left\| p\cdot\theta_{S_p} - p' \right\| + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q] + |labels|$$

initialize plane models θ_0 , θ_1 , θ_2 , ...
from many randomly sampled SIFT matches
in 2 images of the same scene

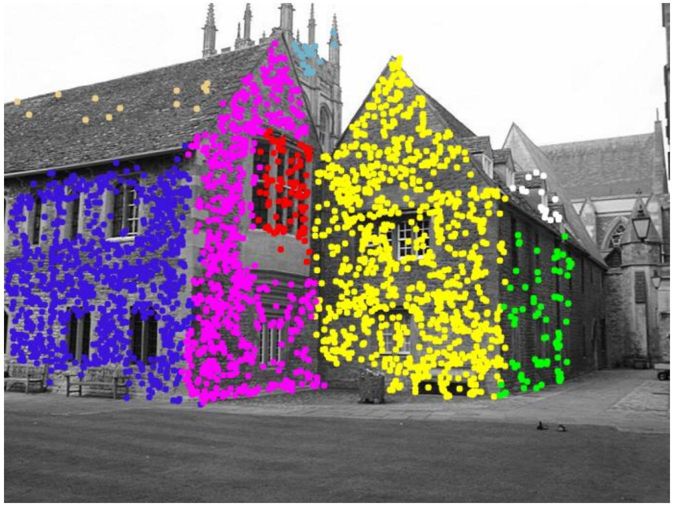

iterate segmentation
and model re-estimation
until convergence

models compete, stable result if sufficiently many

Iterative learning of other models

(could be used for more than 2 labels $s_p \in \{0,1,2,...\}$)

- Geometric multi-model fitting [Isack et al., 2012]

using a-expansion (graph-cuts)

$$E(S,\theta_0,\theta_1,\theta_2...) \quad = \quad \sum_p \left\| p\cdot\theta_{S_p} - p' \right\| \quad + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q] \quad + \quad |\,labels|$$

VIDEO

initialize Fundamental matrices θ_0 , θ_1 , θ_2 , ...
from many randomly sampled SIFT matches
in 2 consecutive frames in video

iterate segmentation
and model re-estimation
until convergence

models compete, stable result if sufficiently many

From color model estimation
to **entropy** and **color consistency**

global optimization in *One Cut*

[Tang et al. ICCV 2013]

Interpretation of log-likelihoods: **entropy** of segment intensities

$$S_i = \{p \in S \mid I_p = i\}$$

pixels of color *i* in *S*

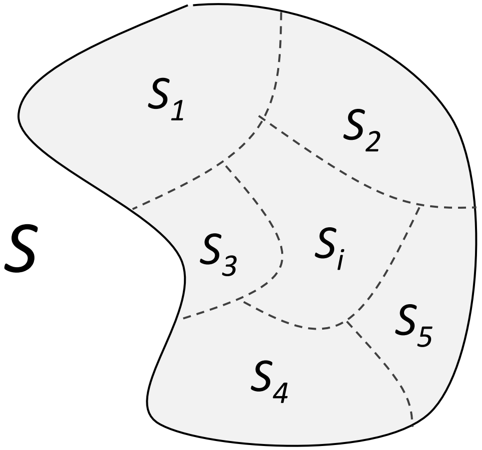

$$-\sum_{p\in S} \ln Pr(I_p \mid \theta)$$

where $\theta = \{p_1, p_2, \dots, p_n\}$

given distribution of intensities

=

$$-\sum_i |S_i| \cdot \ln p_i$$

=

$$-|S| \cdot \sum_i p_i^S \cdot \ln p_i$$

$$p_i^S = \frac{|S_i|}{|S|}$$

probability of intensity i in S

$$p^S = \{p_1^S, p_2^S, \dots, p_n^S\}$$

distribution of intensities observed at S

$$H(S|\theta)$$

cross entropy

of distribution p^S w.r.t. θ

Interpretation of log-likelihoods: **entropy** of segment intensities

joint estimation of S and color models [Rother et al., SIGGRAPH'04, ICCV'09]

$$E(S,\theta_0,\theta_1) = \underbrace{\sum_{p:S_p=0} -\ln \Pr(I_p|\theta_0)}_{|\bar{S}|\cdot\ H(\bar{S}|\theta_0)} + \underbrace{\sum_{p:S_p=1} -\ln \Pr(I_p|\theta_1)}_{|S|\cdot\ H(S|\theta_1)} + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

$$\Big\downarrow \min_{\theta_0,\theta_1}$$

Note: $\overset{\text{cross-entropy}}{H(P|Q)} \geq \overset{\text{entropy}}{H(P)}$ for any two distributions (equality when $Q=P$)

$$E(S) = \overset{\text{entropy of intensities in } \bar{S}}{|\bar{S}|\cdot\ H(\bar{S})} + \overset{\text{entropy of intensities in } S}{|S|\cdot\ H(S)} + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

minimization of segments entropy [Tang et al, ICCV 2013]

Interpretation of log-likelihoods: **entropy** of segment intensities

mixed optimization

[Rother et al., SIGGRAPH’04, ICCV’09]

$$E(S,\theta_0,\theta_1) = \sum_{p:S_p=0} -\ln \Pr(I_p\,|\,\theta_0) + \sum_{p:S_p=1} -\ln \Pr(I_p\,|\,\theta_1) + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

$$|\bar{S}|\cdot\, H(\bar{S}\,|\,\theta_0) \qquad\qquad |S|\cdot\, H(S\,|\,\theta_1)$$

$$\min_{\theta_0,\theta_1}$$

cross-entropy entropy

Note: $H(P|Q) \geq H(P)$ for any two distributions (equality when $Q=P$)

entropy of intensities in $\bar{S}$

entropy of intensities in S

$$E(S) = |\bar{S}|\cdot\, H(\bar{S}) + |S|\cdot\, H(S) + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

binary optimization

[Tang et al, ICCV 2013]

Interpretation of log-likelihoods: **entropy** of segment intensities

$$E(S,\theta_0,\theta_1) = \underbrace{\sum_{p:S_p=0} -\ln \Pr(I_p|\theta_0)}_{|\bar{S}|\cdot\, H(\bar{S}|\theta_0)} + \underbrace{\sum_{p:S_p=1} -\ln \Pr(I_p|\theta_1)}_{|S|\cdot\, H(S|\theta_1)} + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

$$\Big\downarrow \min_{\theta_0,\theta_1}$$

Note: $\overset{\text{cross-entropy}}{H(P|Q)} \geq \overset{\text{entropy}}{H(P)}$ for any two distributions (equality when $Q=P$)

$$E(S) = \boxed{\overset{\text{entropy of intensities in } \bar{S}}{|\bar{S}|\cdot\, H(\bar{S})} + \overset{\text{entropy of intensities in } S}{|S|\cdot\, H(S)}} + \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

common energy for ***categorical clustering***, e.g. [Li *et al*. ICML’04]

Minimizing **entropy** of segments intensities

(intuitive motivation)

$$E(S) = |\bar{S}|\cdot H(\bar{S}) \quad + \quad |S|\cdot H(S) \quad + \sum_{pq\in N} w_{pq}[s_p \neq s_q]$$

break image into two coherent segments
with low entropy of intensities

high entropy segmentation

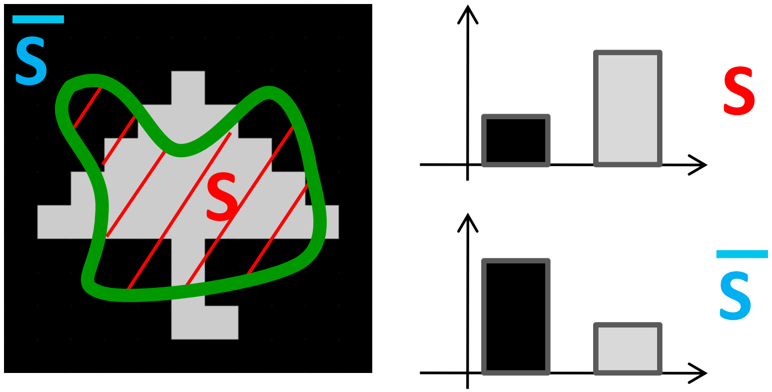

low entropy segmentation

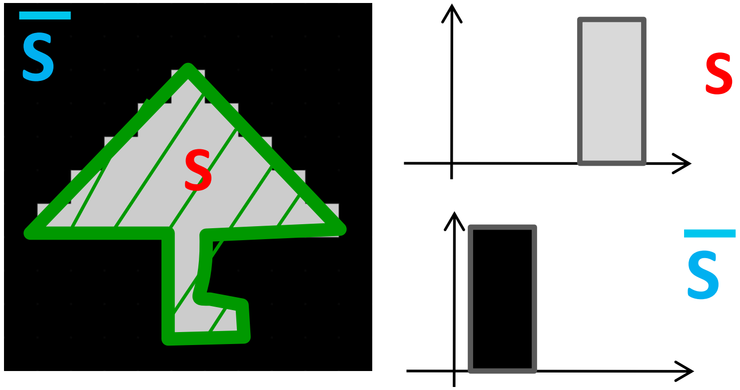

unsupervised image segmentation (like in *Chan-Vese)*

Minimizing **entropy** of segments intensities

(intuitive motivation)

$$E(S) = |\bar{S}| \cdot H(\bar{S}) \quad + \quad |S| \cdot H(S) \quad + \sum_{pq \in N} w_{pq} [s_p \neq s_q]$$

break image into two coherent segments
with low entropy of intensities

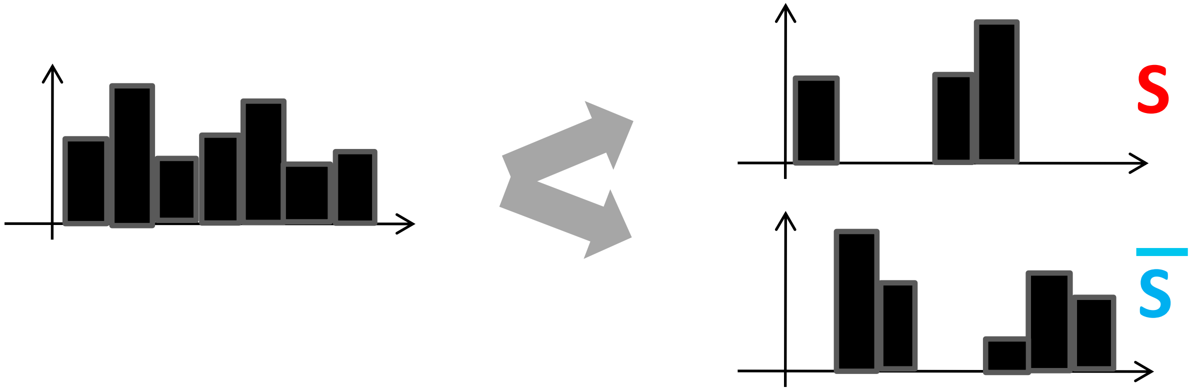

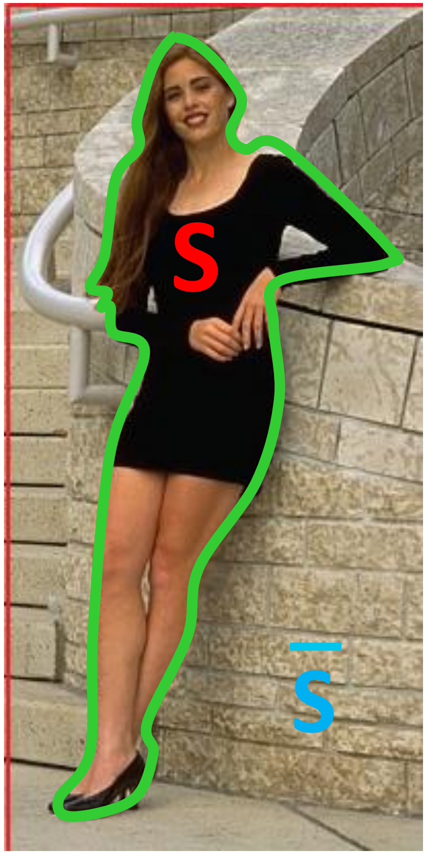

more general than *Chan-Vese* (colors can vary within each segment)

From entropy to color consistency

all pixels $\Omega = \cup \Omega_i$

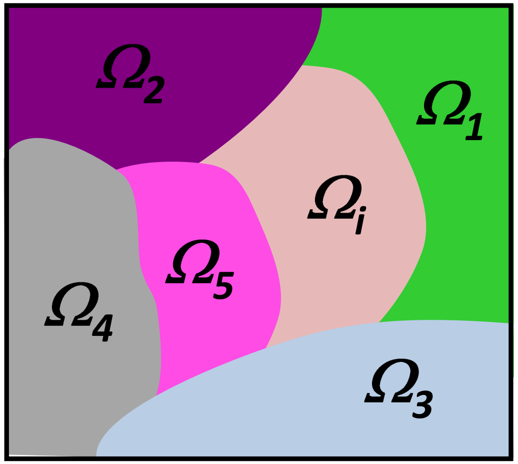

Minimization of entropy encourages pixels Ω_i of the same color bin i to be segmented together

(proof: see next page)

From entropy to color consistency

$$E(S) = \boxed{|\bar{S}|\cdot H(\bar{S}) \quad + |S|\cdot H(S)} + \sum_{pq\in N} w_{pq}[s_p \neq s_q]$$

$$-|\bar{S}|\cdot\sum_i p_i^{\bar{S}} \ln p_i^{\bar{S}} \quad - |S|\cdot\sum_i p_i^{S} \ln p_i^{S}$$

$$-\sum_i |\bar{S}_i|\cdot\ln\tfrac{|\bar{S}_i|}{|\bar{S}|} - \sum_i |S_i|\cdot\ln\tfrac{|S_i|}{|S|}$$

$$\underbrace{\sum_i |\bar{S}_i|}_{|\bar{S}|}\cdot\ln|\bar{S}| - \sum_i |\bar{S}_i|\cdot\ln|\bar{S}_i| \quad + \underbrace{\sum_i |S_i|}_{|S|}\cdot\ln|S| - \sum_i |S_i|\cdot\ln|S_i|$$

$S_i = S \cap \Omega_i$

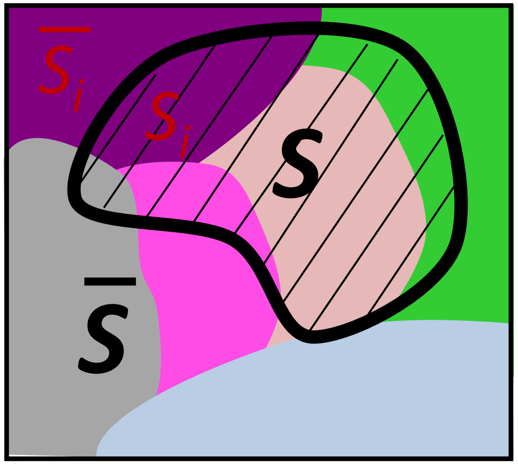

$$\boxed{|\bar{S}|\cdot\ln|\bar{S}| + |S|\cdot\ln|S|} - \sum_i \left(\boxed{|\bar{S}_i|\cdot ln|\bar{S}_i| + |S_i|\cdot ln|S_i|}\right)$$

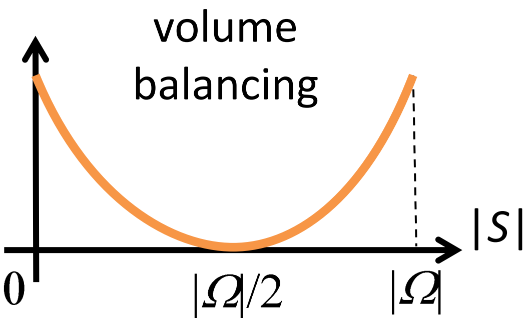

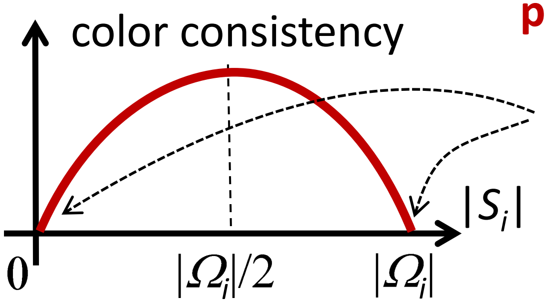

pixels in each color bin *i* prefer to be together (either inside object or background)

From entropy
to color consistency

segmentation *S*
with better
color consistency

$S_i = S \cap \Omega_i$

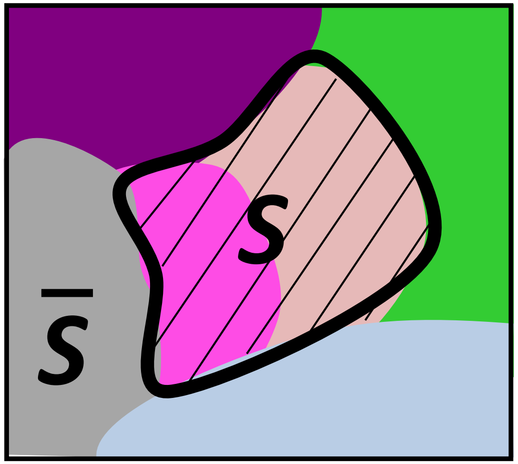

$$\boxed{|\bar{S}|\cdot\ln|\bar{S}| + |S|\cdot\ln|S|} - \sum_i (\boxed{|\bar{S}_i|\cdot ln|\bar{S}_i| + |S_i|\cdot ln|S_i|})$$

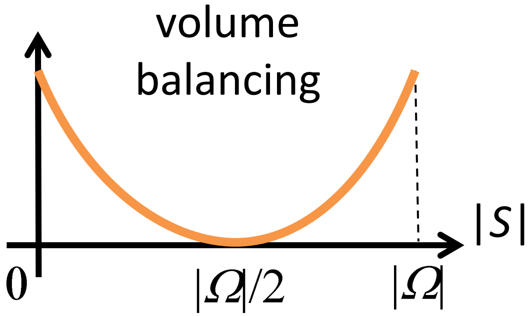

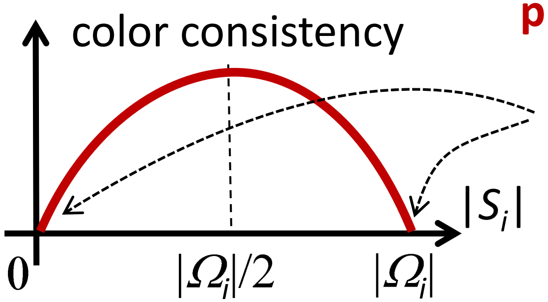

pixels in each color bin *i*
prefer to be together
(either inside object
or background)

From entropy
to color consistency

In many applications, this term can be either dropped or replaced with simple unary ballooning [Tang et al. ICCV 2013]

Graph-cut constructions for similar cardinality terms (for superpixel consistency) [Kohli et al. IJCV'09]

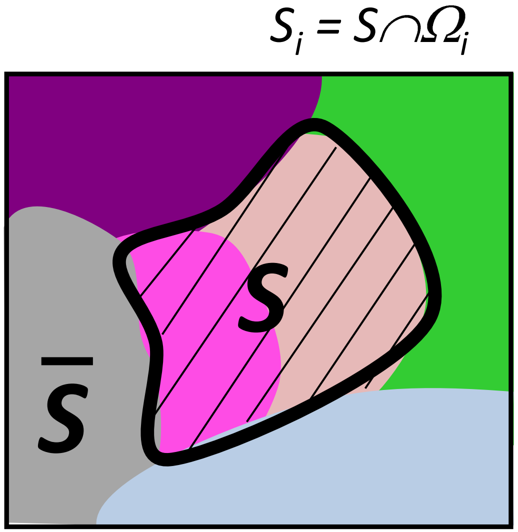

~~**convex function of cardinality $|S|$**~~
~~***(non-submodular)***~~

concave function of cardinality $|S_i|$
(submodular)

$$\boxed{\xcancel{|\bar{S}|\cdot\ln|\bar{S}|+|S|\cdot\ln|S|}} - \sum_i \left(\boxed{|\bar{S}_i|\cdot ln|\bar{S}_i|+|S_i|\cdot ln|S_i|}\right)$$

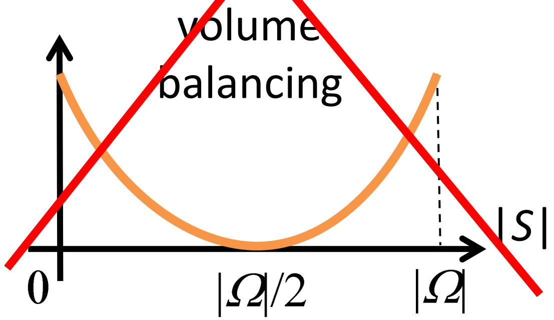

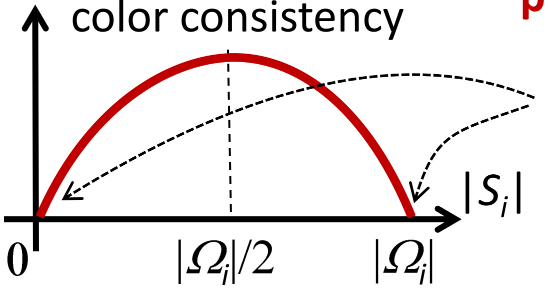

pixels in each color bin *i* prefer to be together (either inside object or background)

From entropy to color consistency

In many applications, this term can be either dropped or replaced with simple unary ballooning [Tang et al. ICCV 2013]

L_1 color separation works better in practice [Tang et al. ICCV 2013] (also, simpler construction)

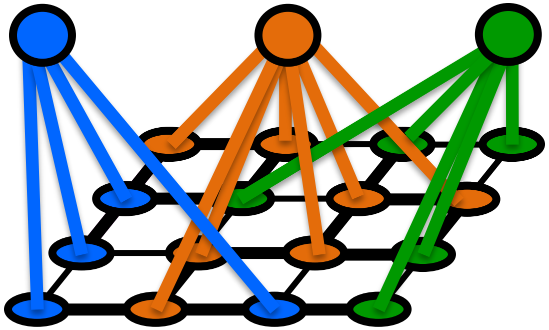

connect pixels in each color bin to corresponding auxiliary nodes

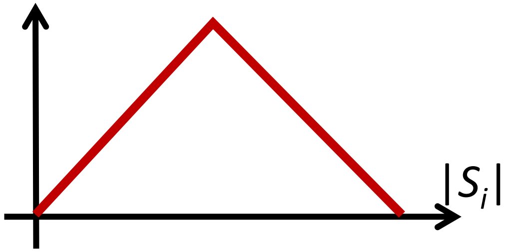

~~**convex function of cardinality |S|**~~

~~***(non-submodular)***~~

$$\cancel{\boxed{|\bar{S}|\cdot\ln|\bar{S}|+|S|\cdot\ln|S|}} - \sum_i (\boxed{|\bar{S}_i|\cdot ln|\bar{S}_i|+|S_i|\cdot ln|S_i|}) + \sum_{pq\in N} w_{pq}[s_p \neq s_q]$$

~~volume balancing~~

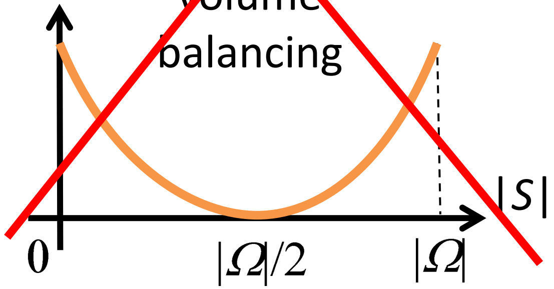

color consistency

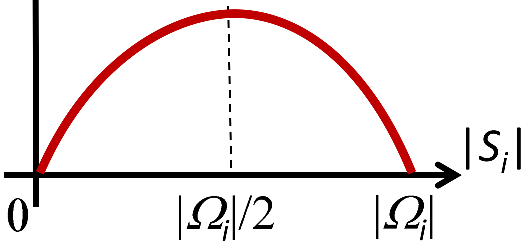

boundary smoothness

smoothness + color consistency

- One Cut

  [Tang, et al., ICCV'13]

  guaranteed global minimum

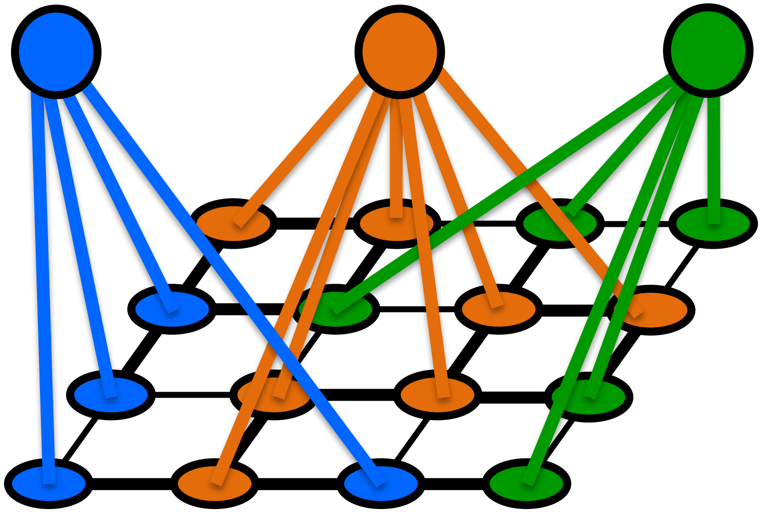

connect pixels in each color bin
to corresponding auxiliary nodes

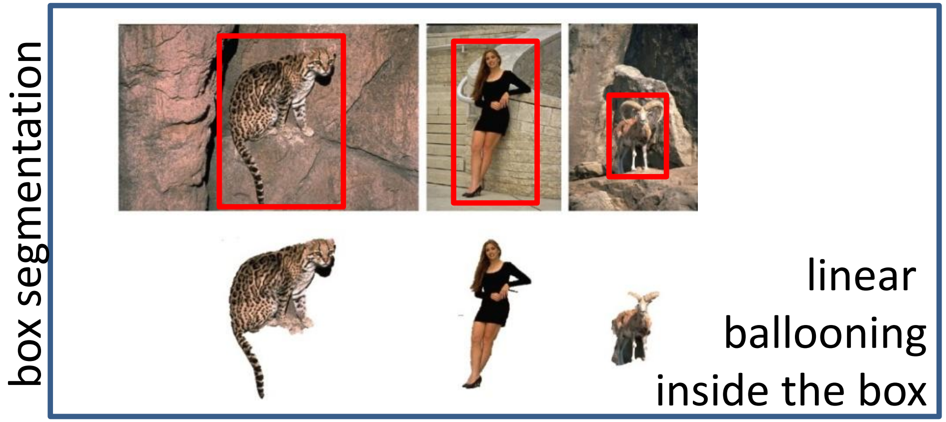

linear
ballooning
inside the box

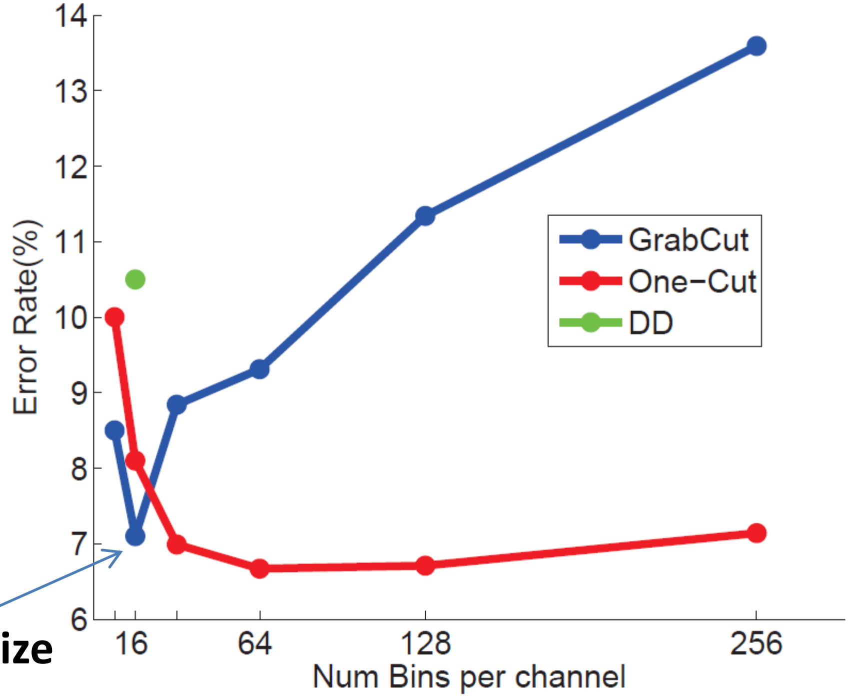

Grabcut is
sensitive to **bin size**

smoothness + color consistency

- One Cut

  [Tang, et al., ICCV’13]

  guaranteed global minimum

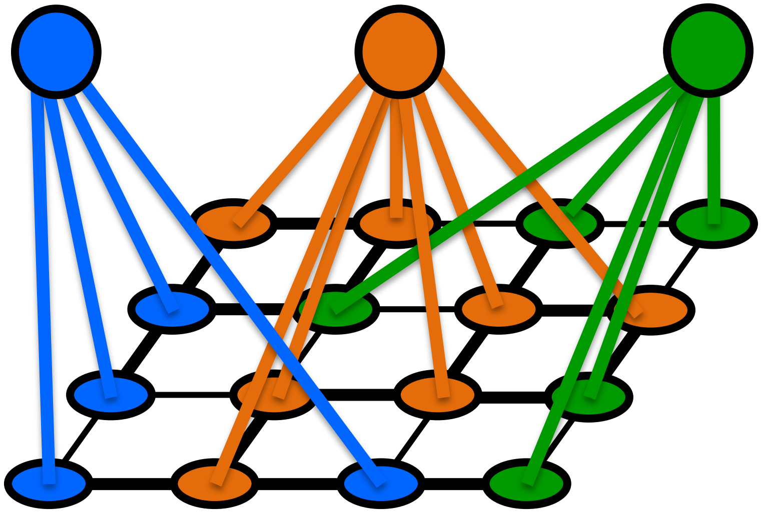

connect pixels in each color bin
to corresponding auxiliary nodes

box segmentation

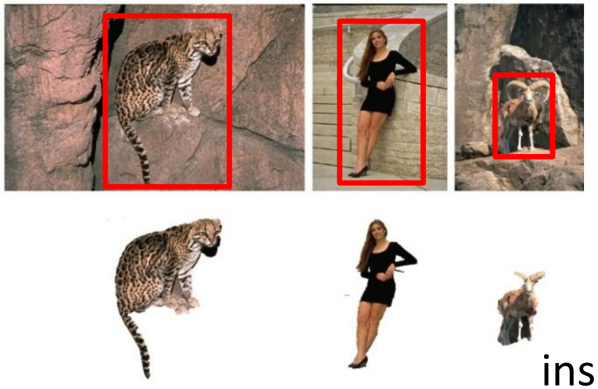

linear
ballooning
inside the box

from seeds

ballooning from
hard constraints

saliency-based
segmentation

linear
ballooning from
saliency measure

photo-consistency + smoothness + color consistency

- Color consistency can be integrated into binary stereo

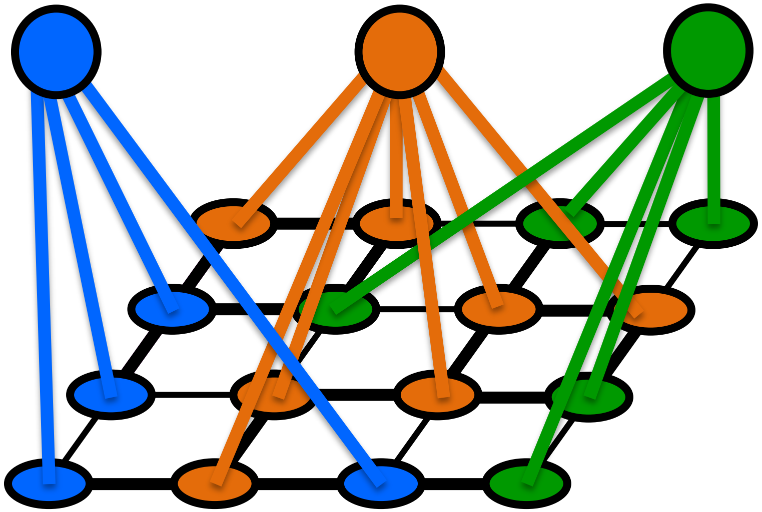

connect pixels in each color bin
to corresponding auxiliary nodes

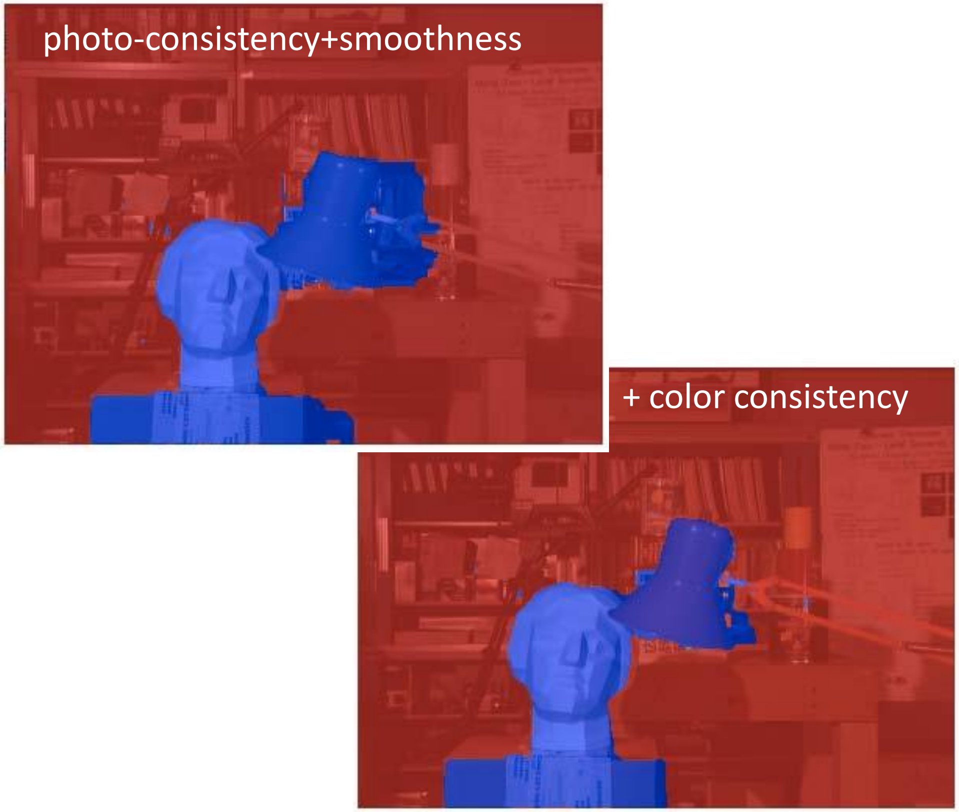

Approximating:

- Convex cardinality potentials
- Distribution consistency
- Other high-order region terms

General Trust Region Approach (overview)

$$\mathbf{E(S)} = \mathrm{H}\mathbf{(S)} + \mathrm{B}\mathbf{(S)}$$

hard

submodular (easy)

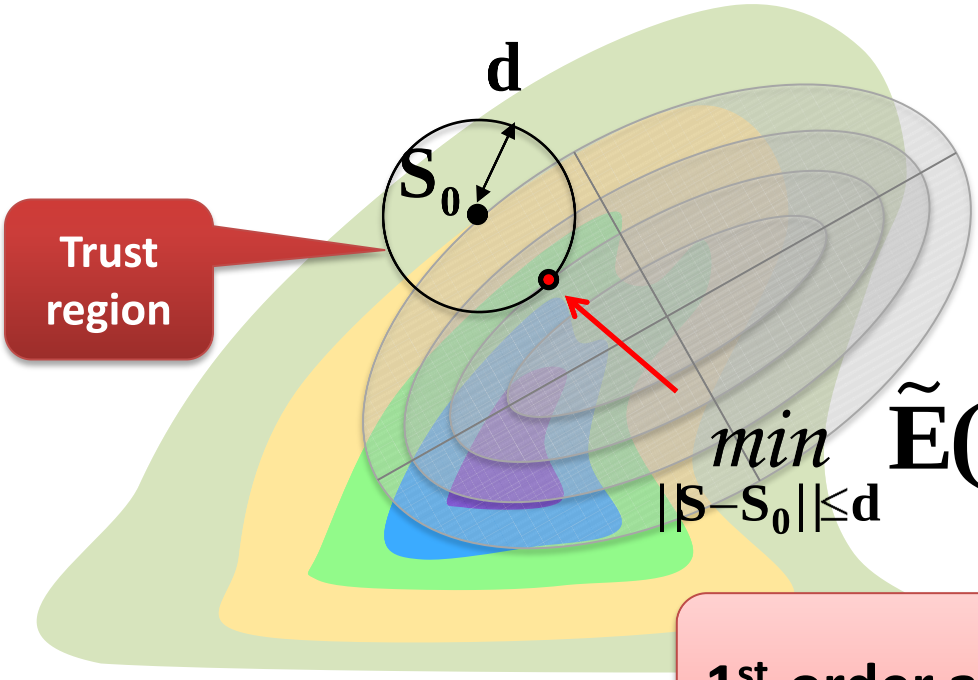

Trust region

$$\min_{||\mathbf{S}-\mathbf{S}_0||\le \mathbf{d}} \tilde{\mathbf{E}}\mathbf{(S)} = \mathrm{U}_0\mathbf{(S)} + \mathrm{B}\mathbf{(S)}$$

1st-order approximation for H(S)

General Trust Region Approach (overview)

- Constrained optimization

minimize $\tilde{\mathbf{E}}(\mathbf{S}) = \mathbf{U}_0(\mathbf{S}) + \mathbf{B}(\mathbf{S})$

$\mathbf{s.t.} \| \mathbf{S} - \mathbf{S}_0 \| \leq \mathbf{d}$

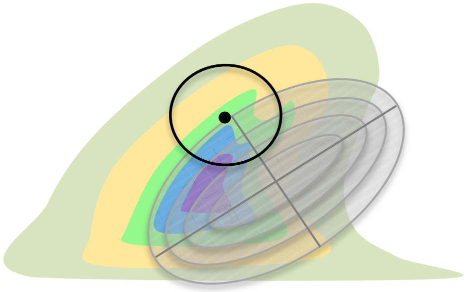

- Unconstrained Lagrangian Formulation

minimize $\mathbf{L}_\lambda(\mathbf{S}) = \mathbf{U}_0(\mathbf{S}) + \mathbf{B}(\mathbf{S}) + \lambda \| \mathbf{S} - \mathbf{S}_0 \|$

can be approximated with unary terms
[Boykov,Kolmogorov,Cremers,Delong, ECCV'06]

Approximating L_2 distance $\| S - S_0 \|$

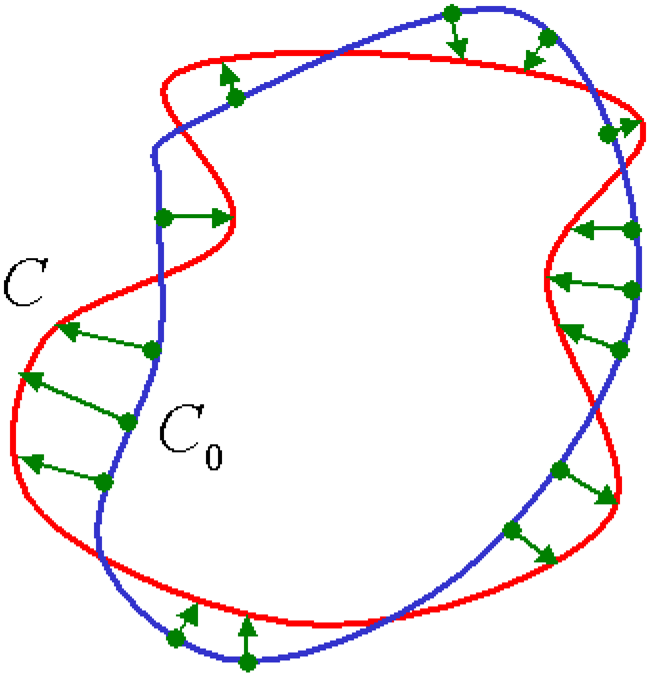

d_p - signed distance map from C_0

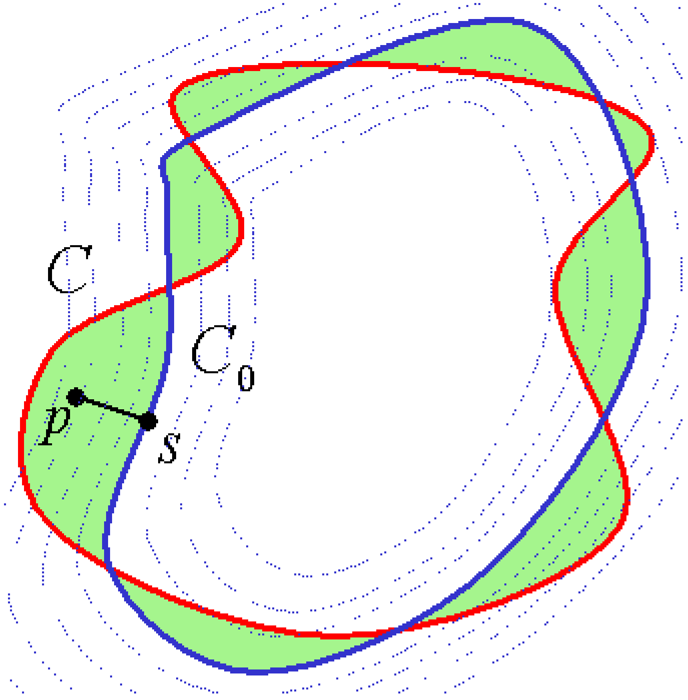

$$\langle dC, dC \rangle = \int_{C_0} dC_s^2 \cdot ds \qquad \approx 2 \int_{\Delta C} d_p \cdot dp \qquad \approx 2 \sum_p d_p \cdot (s_p - s_p^o)$$

unary potentials [Boykov et al. ECCV 2006]

Trust Region Approximation

non-submodular term

Linear approx. at S_0

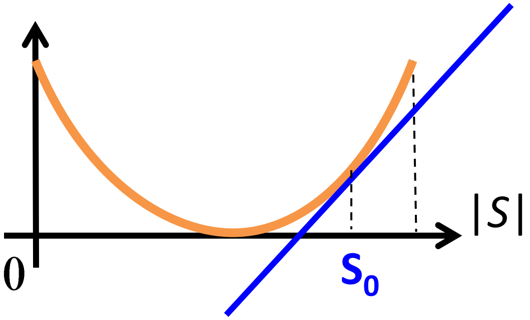

volume constraint

submodular terms

$$+ \sum_p -ln\Pr(I_p \,|\, \theta_{S_p}) \;+\; \sum_{pq\in N} w_{pq}\cdot[s_p \neq s_q]$$

appearance log-likelihoods

boundary length

L_2 distance to S_0

$$+ \quad \lambda\cdot\sum_p d_p(s_p - s_p^o)$$

trust region

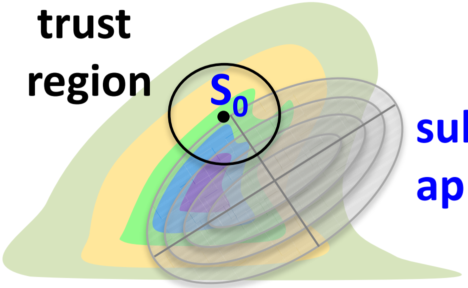

submodular approx.

Volume Constraint for Vertebrae segmentation

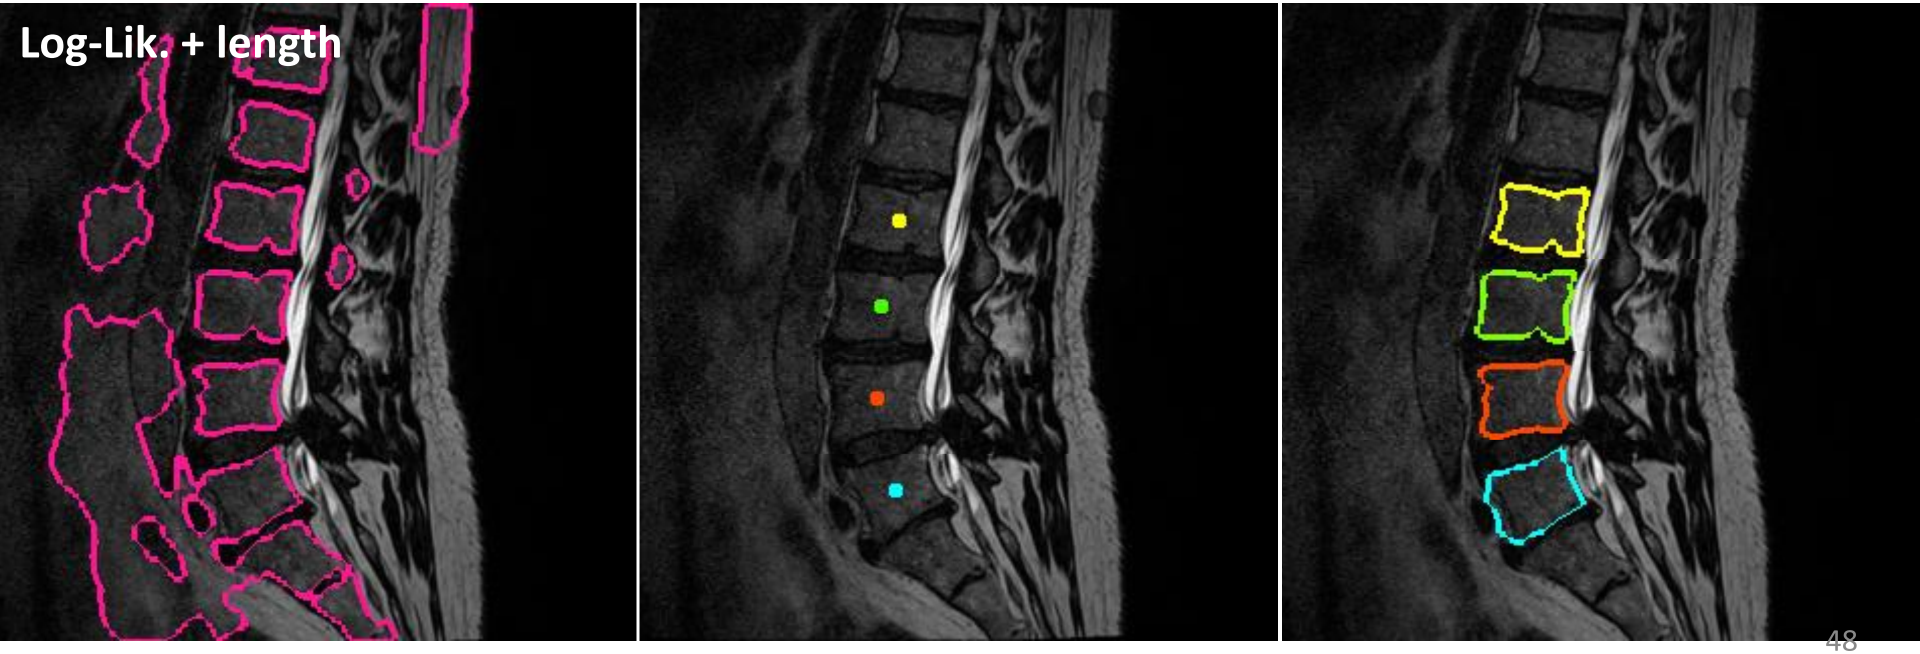

Back to entropy-based segmentation

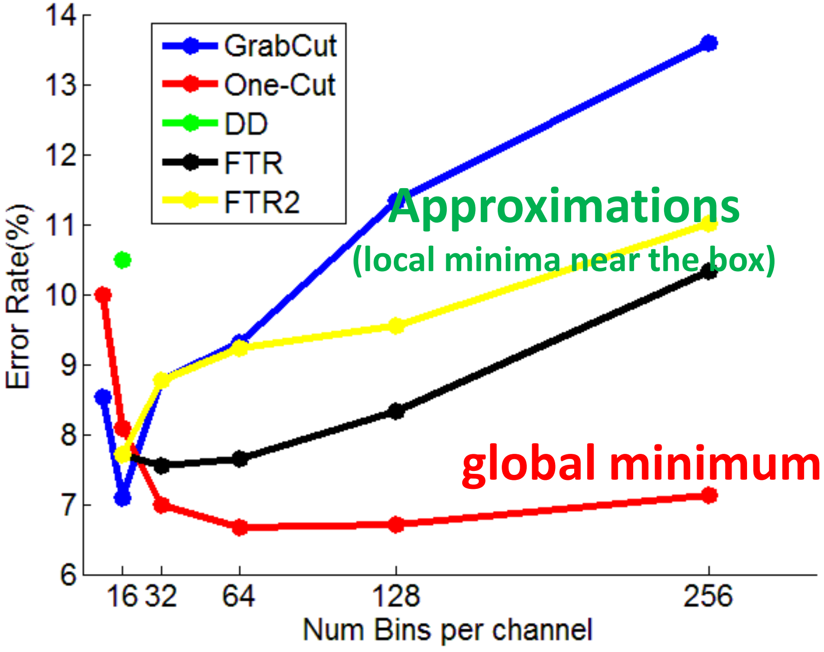

Approximations

(local minima near the box)

global minimum

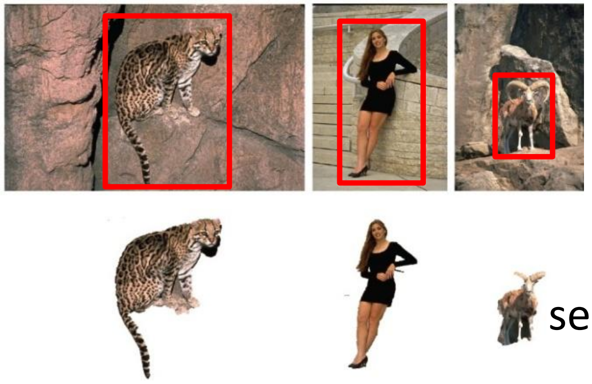

Interactive segmentation with box

non-submodular term

volume balancing

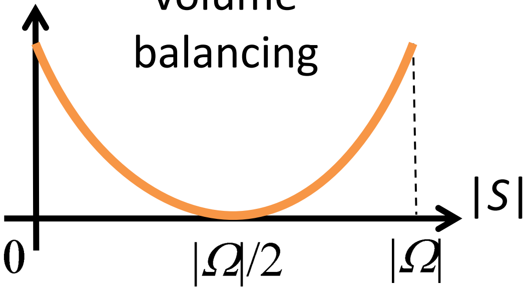

+

submodular terms

color consistency

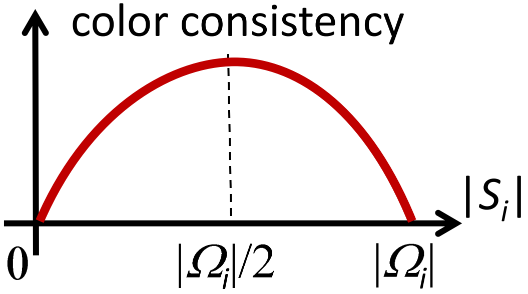

boundary smoothness

$$+ \sum_{pq \in N} w_{pq}[s_p \neq s_q]$$

Trust Region Approximation

Surprisingly, TR outperforms QPBO, DD, TRWS, BP, etc. on many high-order [CVPR'13] and/or non-submodular problems [arXiv13]

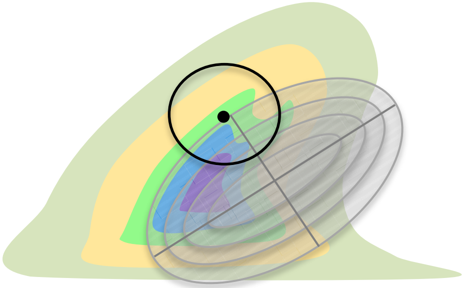

Curvature

Pair-wise smoothness: limitations

- discrete metrication errors

    - resolved by higher connectivity
    - continuous convex formulations

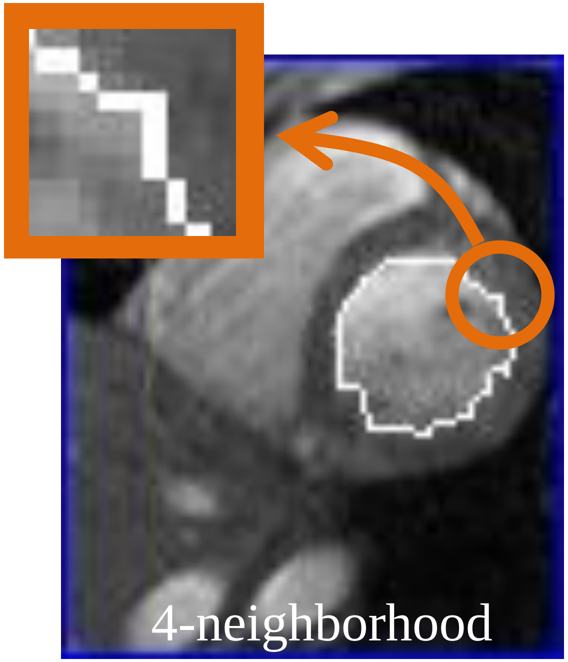

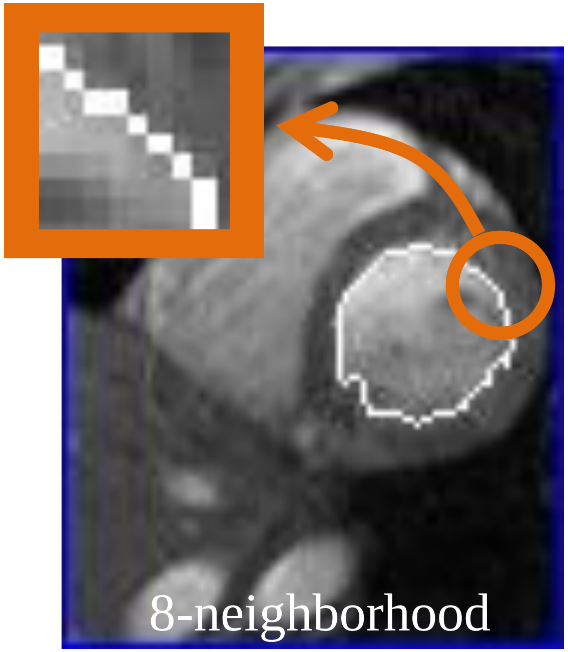

Pair-wise smoothness: limitations

- boundary over-smoothing (a.k.a. *shrinking bias*)

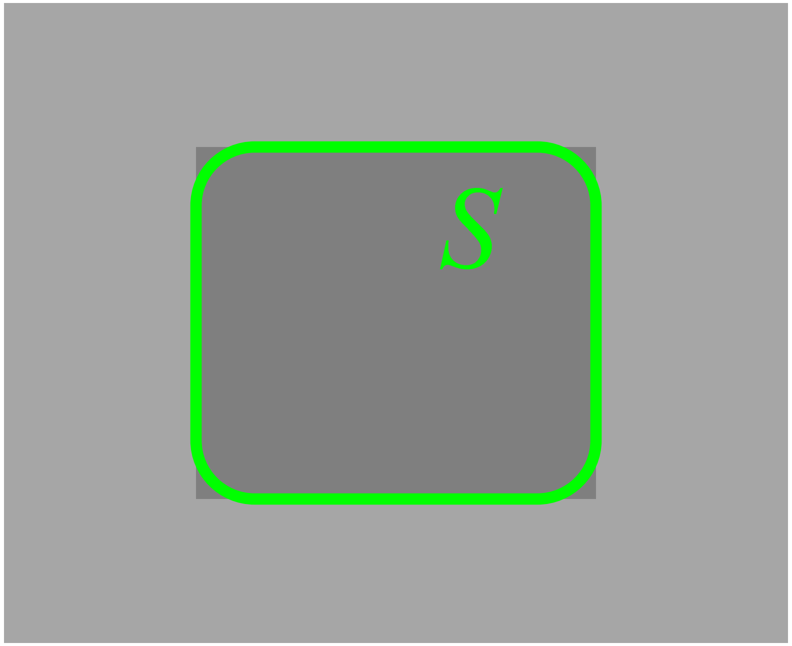

Pair-wise smoothness: limitations

- boundary over-smoothing (a.k.a. *shrinking bias*)
  - needs higher-order smoothness
  - curvature

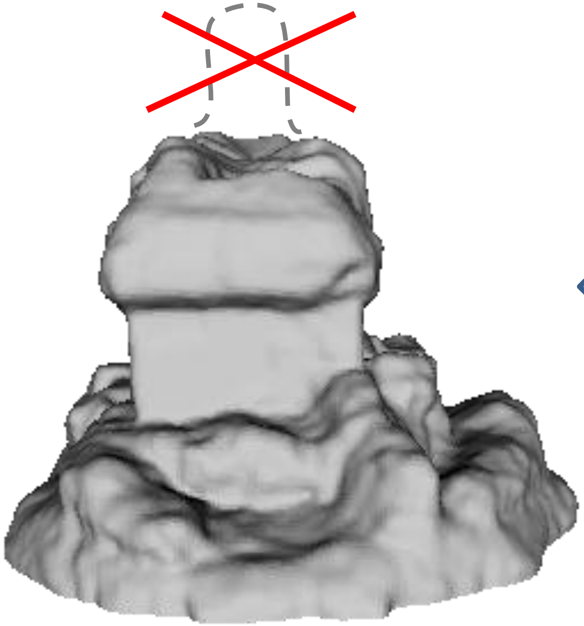

multi-view reconstruction
[Vogiatzis et al. 2005]

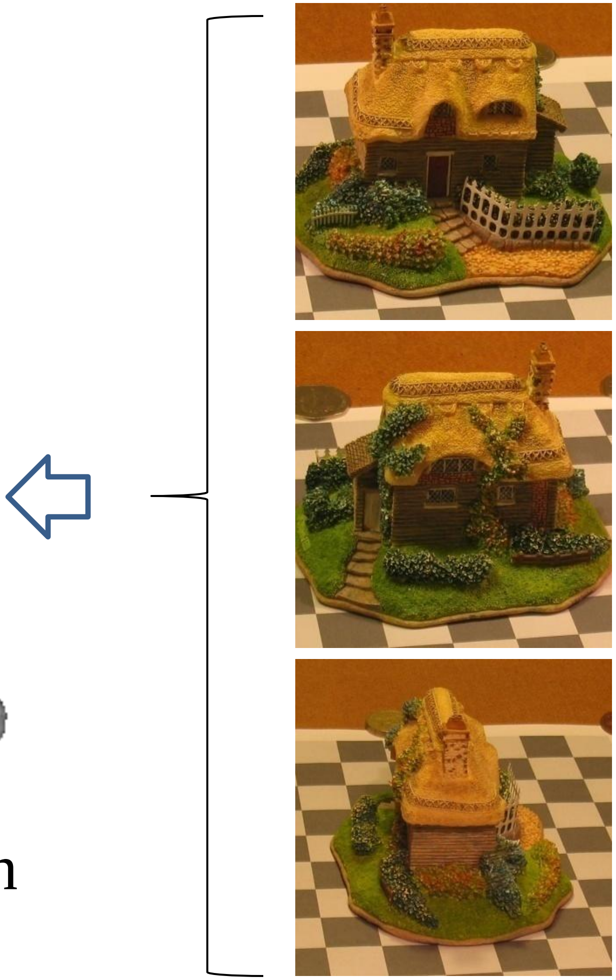

Higher-order smoothness & curvature for discrete regularization

- Geman and Geman 1983 (line process, simulated annealing)

- Second-order stereo and surface reconstruction
  - Li & Zuker 2010 (loopy belief propagation)
  - Woodford et al. 2009 (fusion of proposals, QPBO)
  - Olsson et al. 2012-13 (fusion of planes, nearly submodular)

- Curvature in segmentation:
  - Schoenemann et al. 2009 (complex, LP relaxation, many extra variables)
  - Strandmark & Kahl 2011 (complex, LP relaxation,…)
  - El-Zehiry & Grady 2010 (grid, 3-clique, only 90 degree accurate, QPBO)
  - Shekhovtsov et al. 2012 (grid patches, approximately learned, QPBO)
  - Olsson et al. 2013 (grid patches, integral geometry, partial enumeration)
  - Nieuwenhuis et al 2014? (grid, 3-cliques, integral geometry, trust region)

this talk

good approximation of curvature, better and faster optimization

practical !

the rest of the talk:

- **Absolute curvature** regularization on a grid

  [Olsson, Ulen, Boykov, Kolmogorov - ICCV 2013]

- **Squared curvature** regularization on a grid

  [Nieuwenhuis, Toppe, Gorelick, Veksler, Boykov - arXiv 2013]

Absolute Curvature

$$\oint_{\partial S} |\kappa| \cdot ds$$

Motivating example: for any convex shape $\oint_{\partial S} |\kappa| \cdot ds = 2\pi$

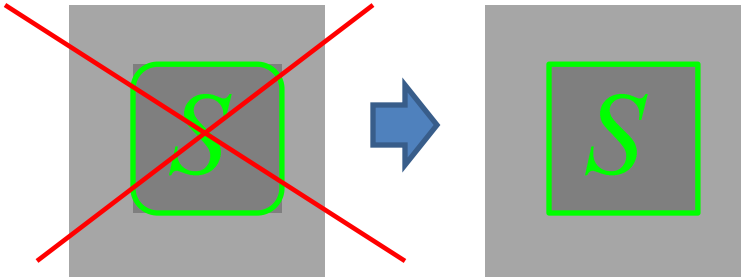

- no shrinking bias
- thin structures

Absolute Curvature

$$\oint_{\partial S} |\kappa| \cdot ds \approx \sum_n \theta_n$$

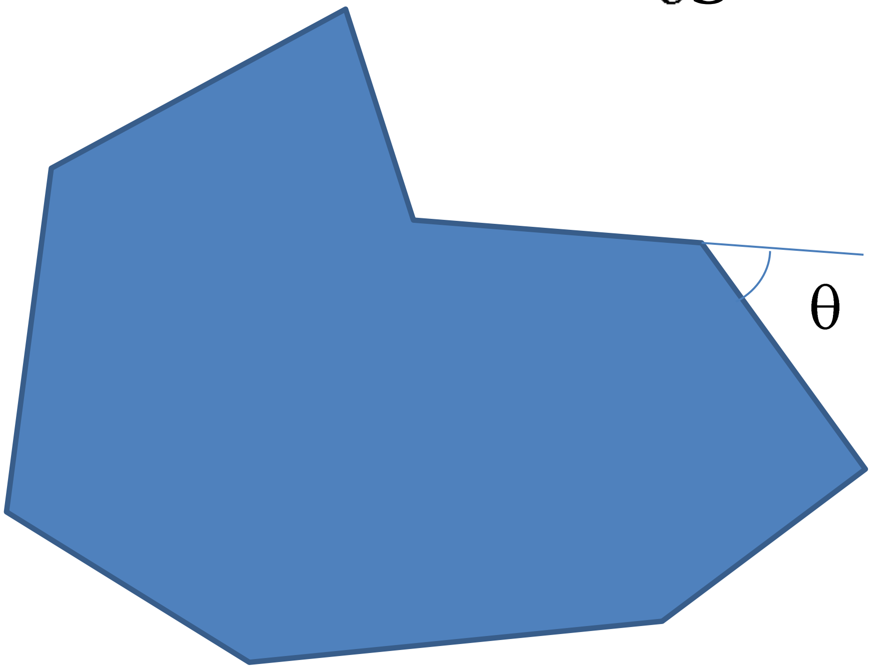

easy to estimate
via approximating
polygons

polygons also work for $|\kappa|^p$
[Bruckstein et al. 2001]

curvature on a cell complex (standard geometry)

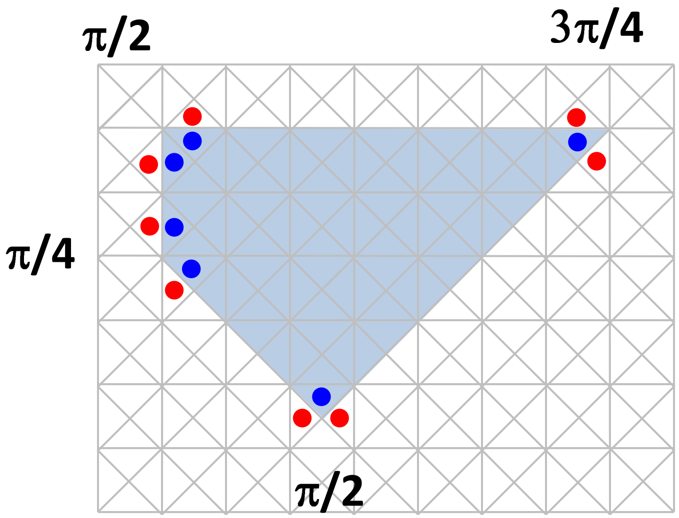

4- or 3-cliques on a cell complex

- Schoenemann et al. 2009
- Strandmark & Kahl 2011

solved via LP relaxations

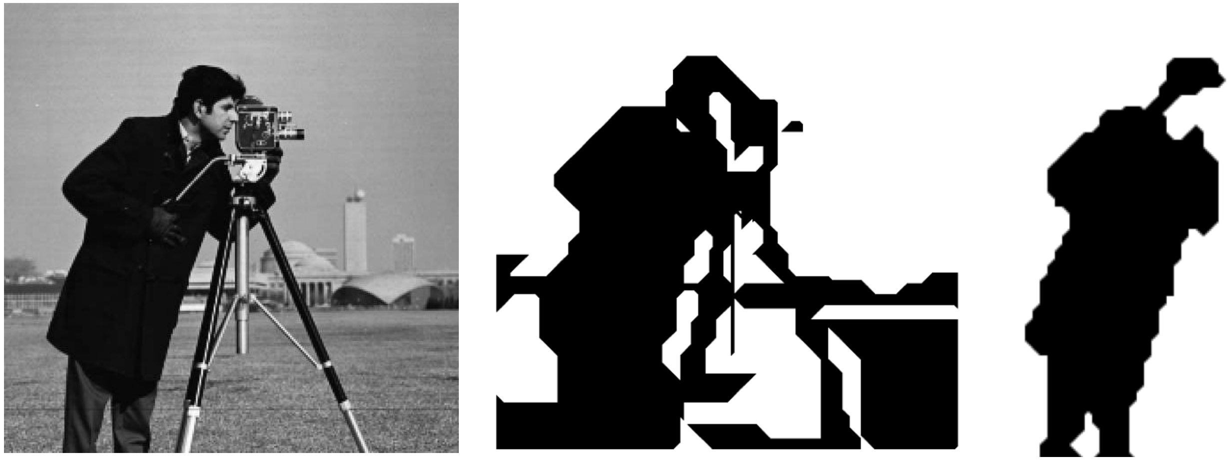

curvature on a cell complex
(standard geometry)

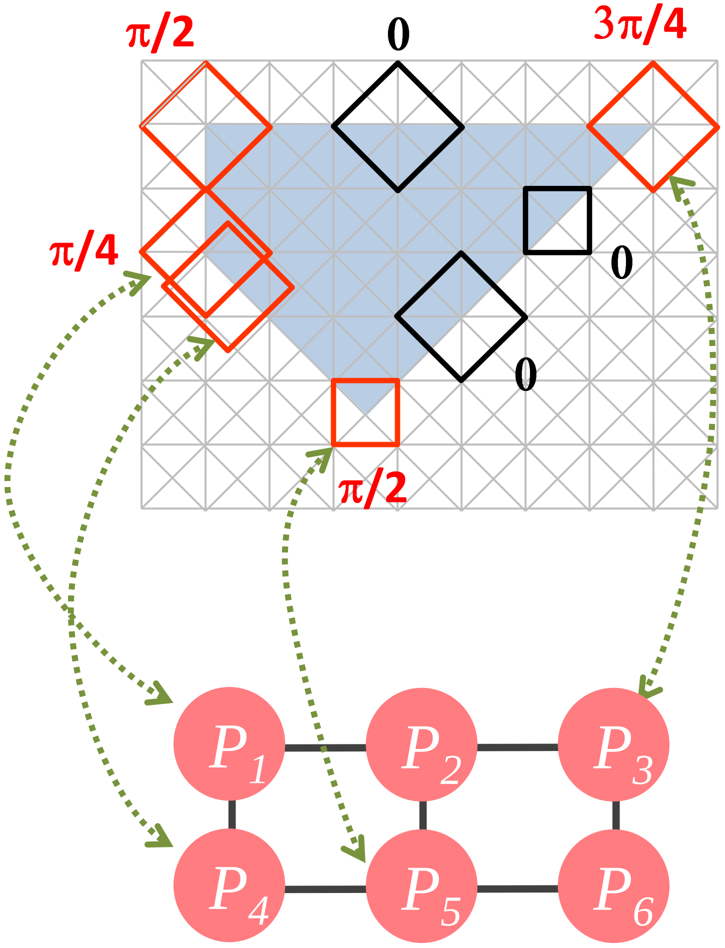

reduction to pair-wise
Constrain Satisfaction Problem

zero gap

cell-patch cliques on a complex

- Olsson et al., ICCV 2013

partial enumeration + TRWS

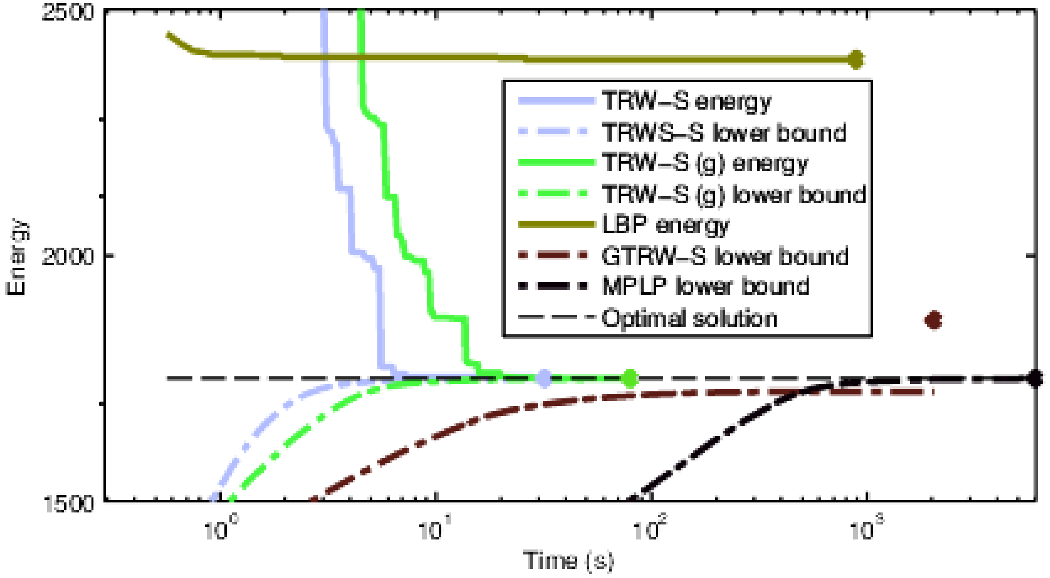

curvature on a cell complex (standard geometry)

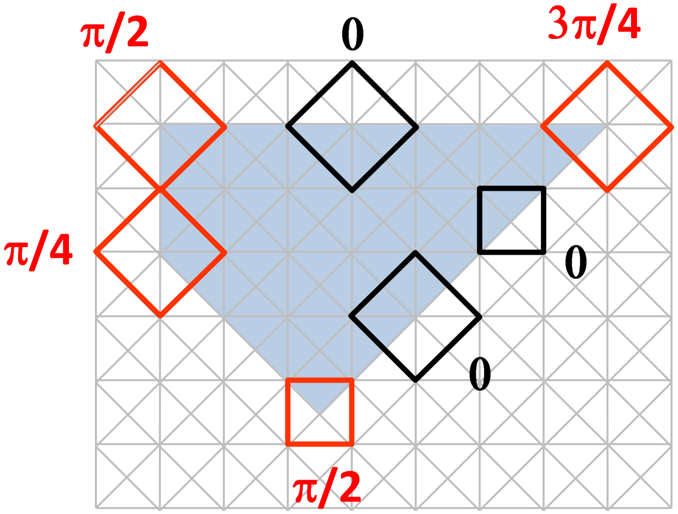

representative cell-patches

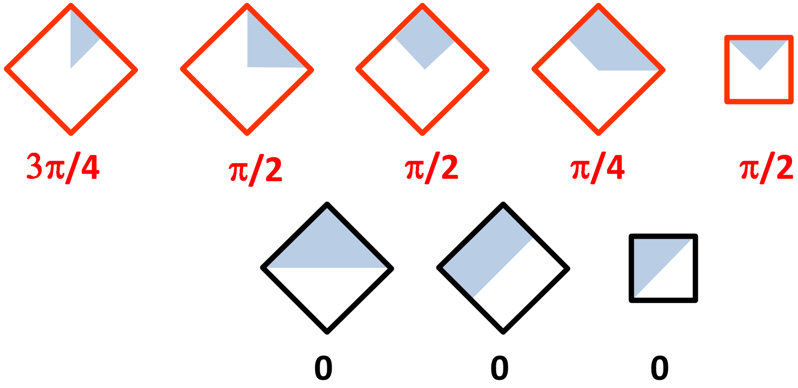

curvature on a pixel grid (integral geometry)

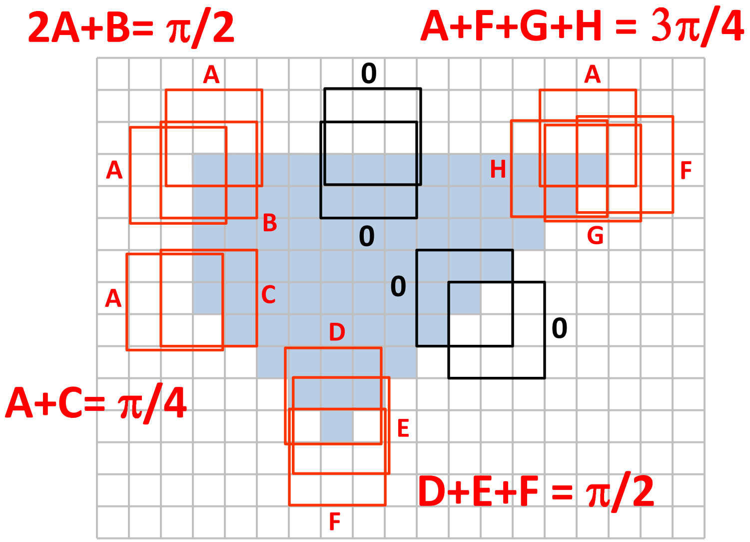

representative pixel-patches

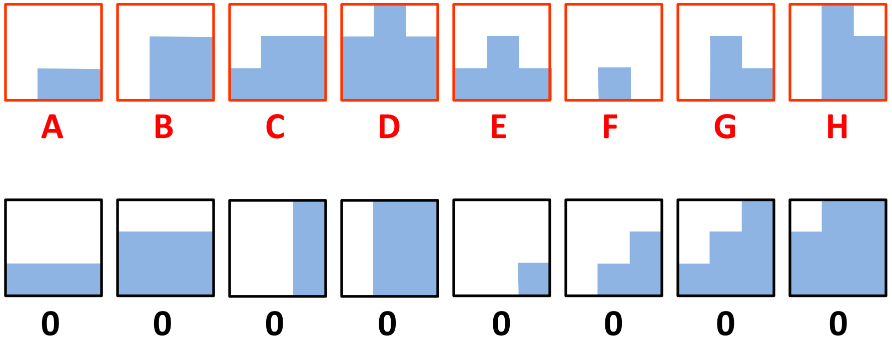

integral approach to absolute curvature on a grid

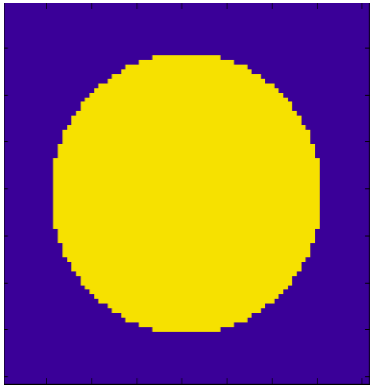

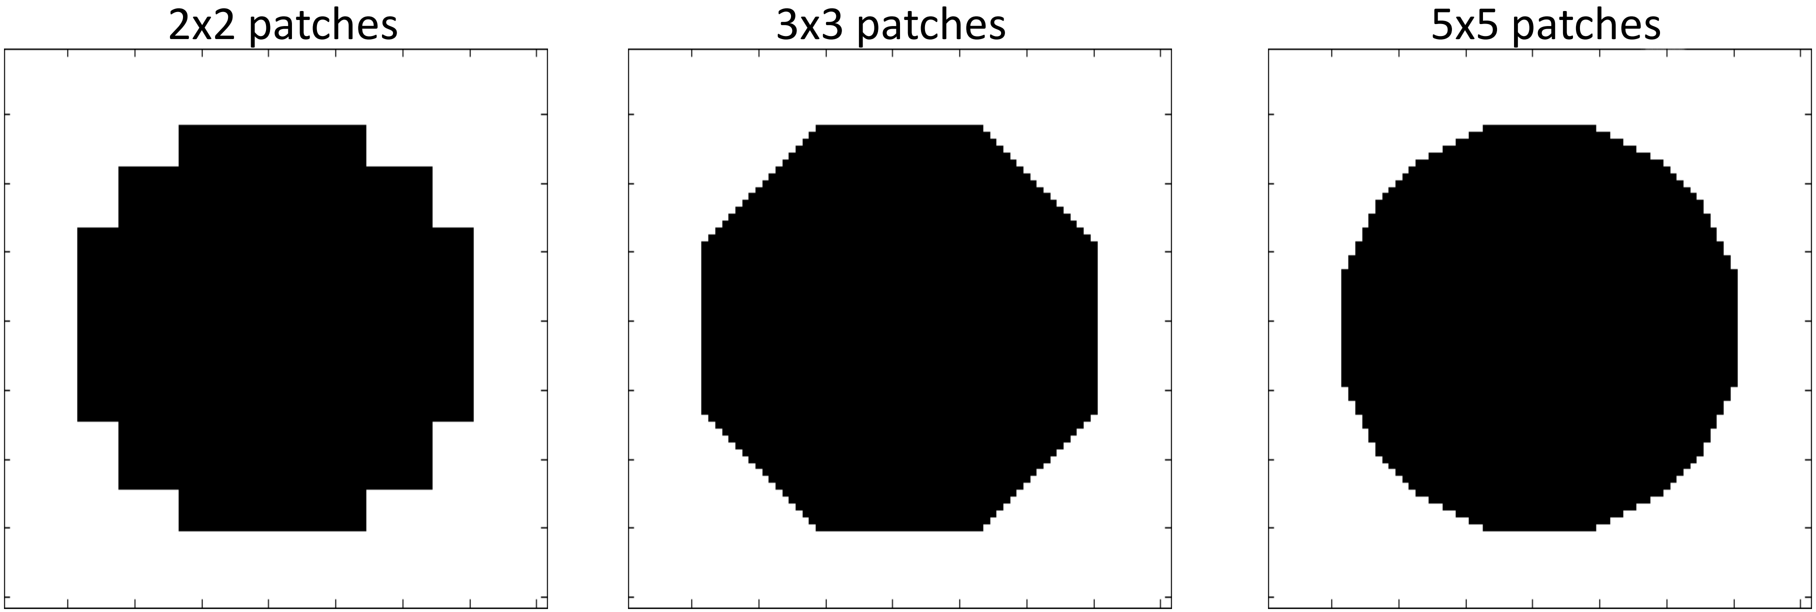

zero gap

integral approach to absolute curvature on a grid

2x2 patches

3x3 patches

5x5 patches

zero gap

Squared Curvature with 3-cliques

$$\oint_{\partial S} \kappa^2 \cdot ds$$

Nieuwenhuis et al., arXiv 2013

general intuition example

3-cliques

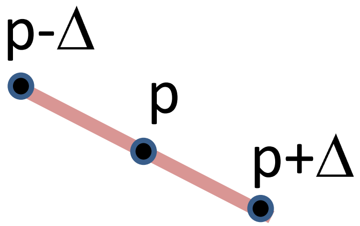

with configurations (0,1,0) and (1,0,1)

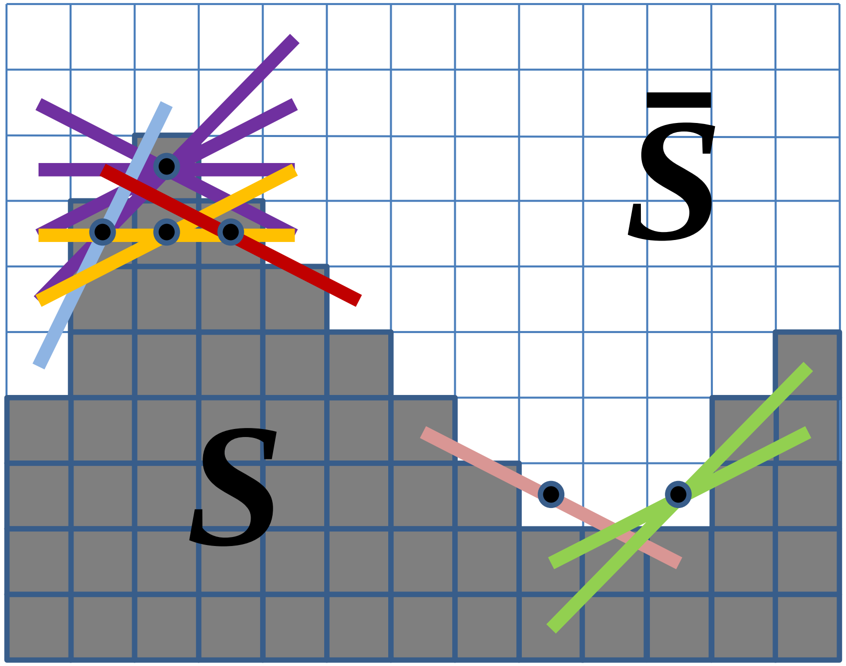

more responses where curvature is higher

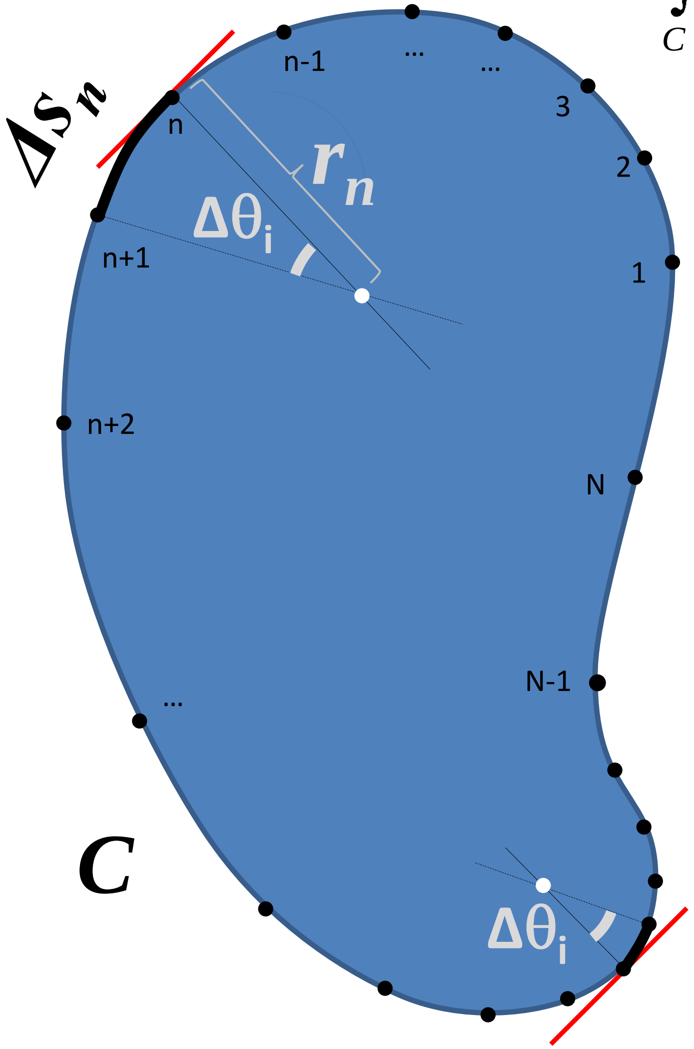

$$\int_C \kappa^2(s)\cdot ds \approx \sum_{n=1}^{N} \kappa_n^2 \cdot \Delta s_n \approx \sum_{n=1}^{N} |\kappa_n|\cdot \Delta\theta_{i(n)}$$

$$|\kappa_n| \cdot \Delta s_n \approx \Delta\theta_{i(n)}$$

5x5 neighborhood

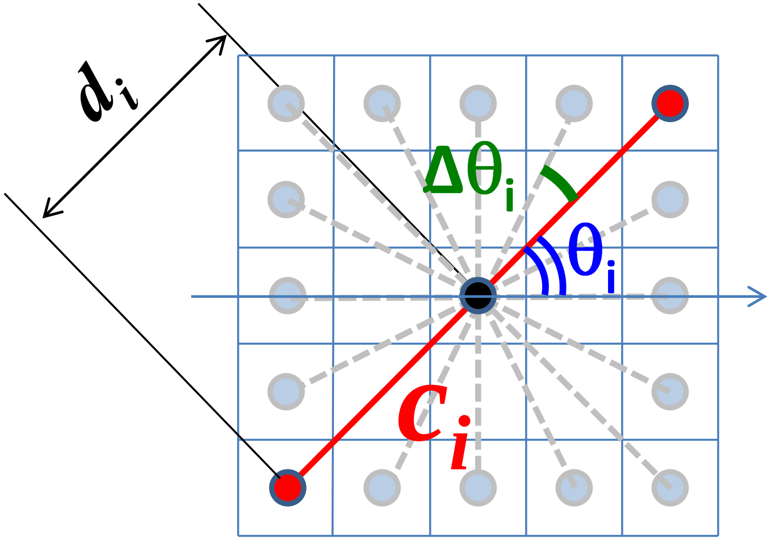

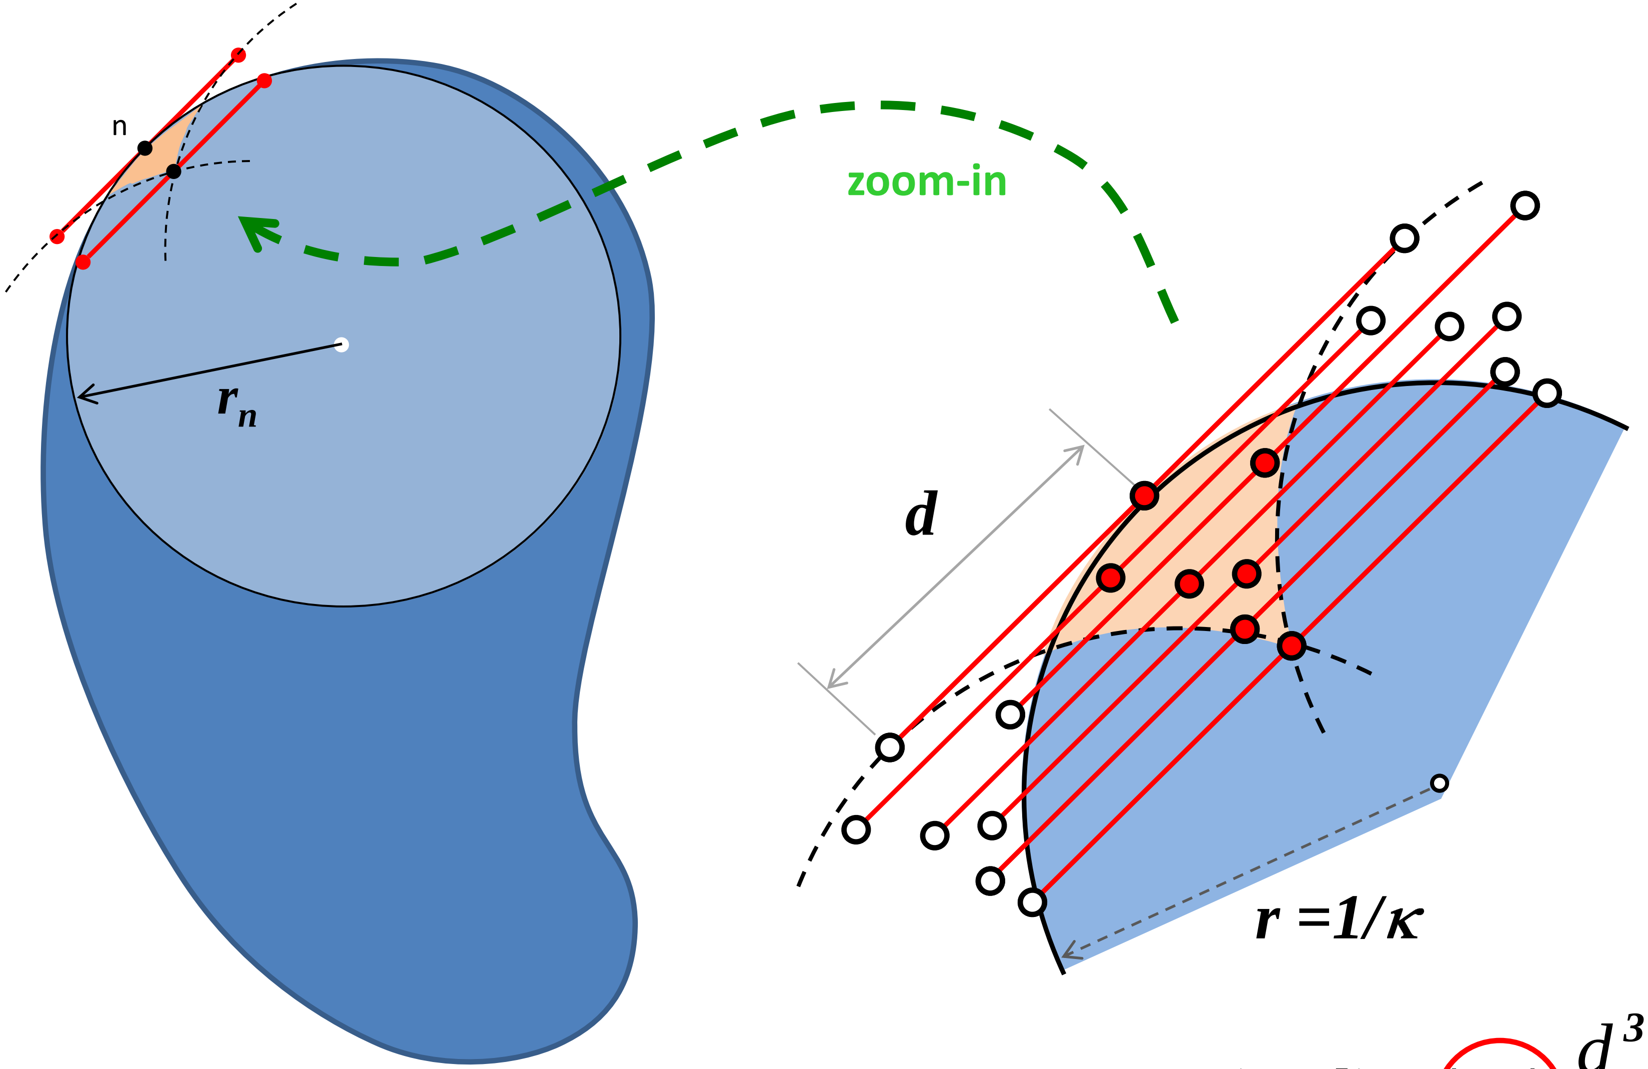

$$R(\kappa,d) \approx \boxed{|\kappa|} \cdot \frac{d^3}{4}$$

Thus, appropriately weighted 3-cliques estimate squared curvature integral

Experimental evaluation

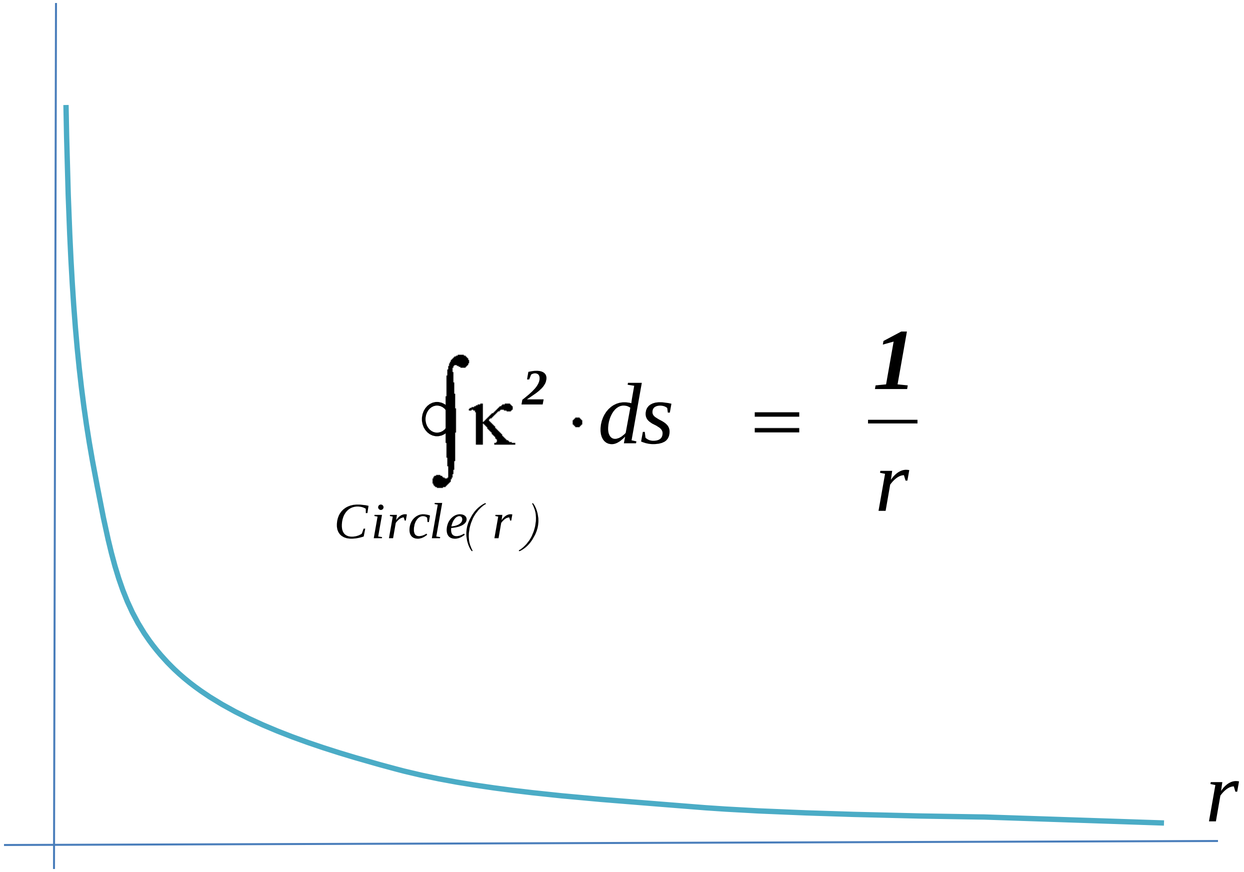

$$\oint\limits_{Circle(r)} \kappa^2 \cdot ds = \frac{1}{r}$$

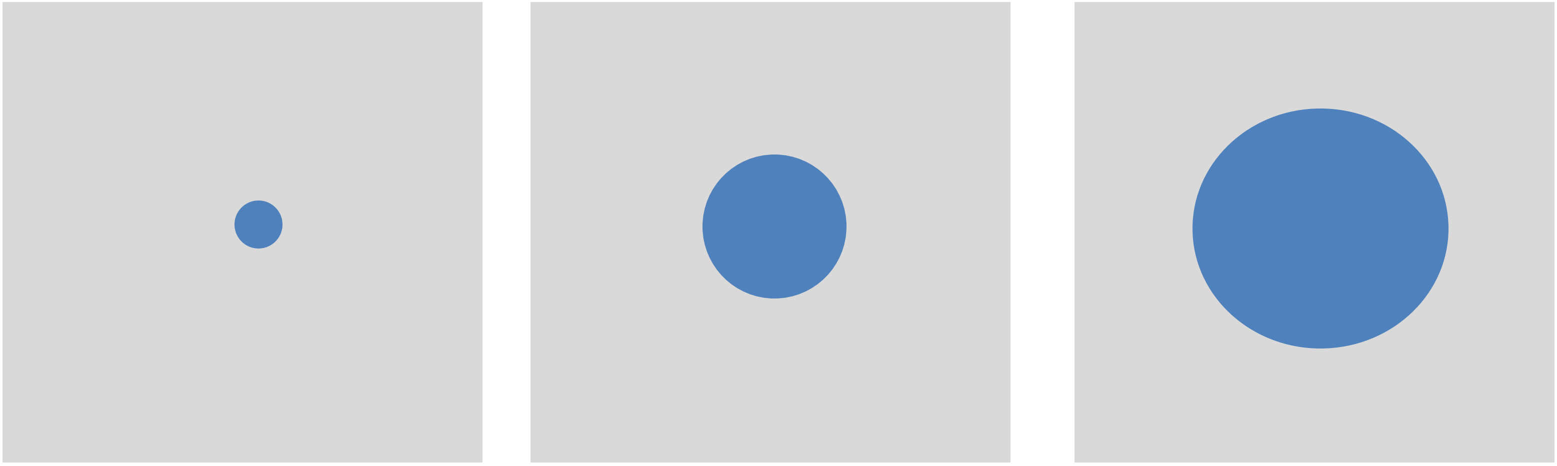

Experimental evaluation

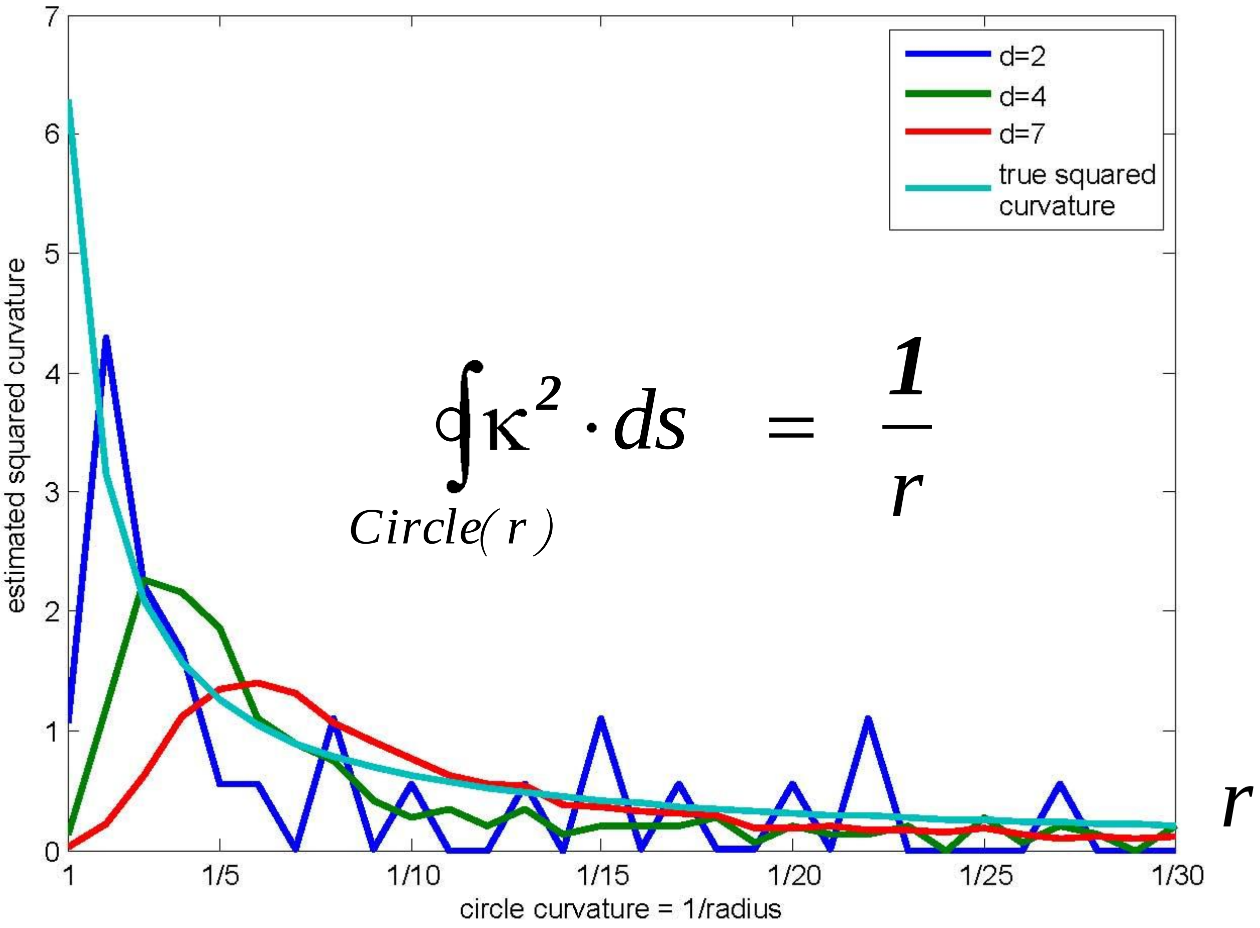

$$\oint_{Circle(r)} \kappa^2 \cdot ds = \frac{1}{r}$$

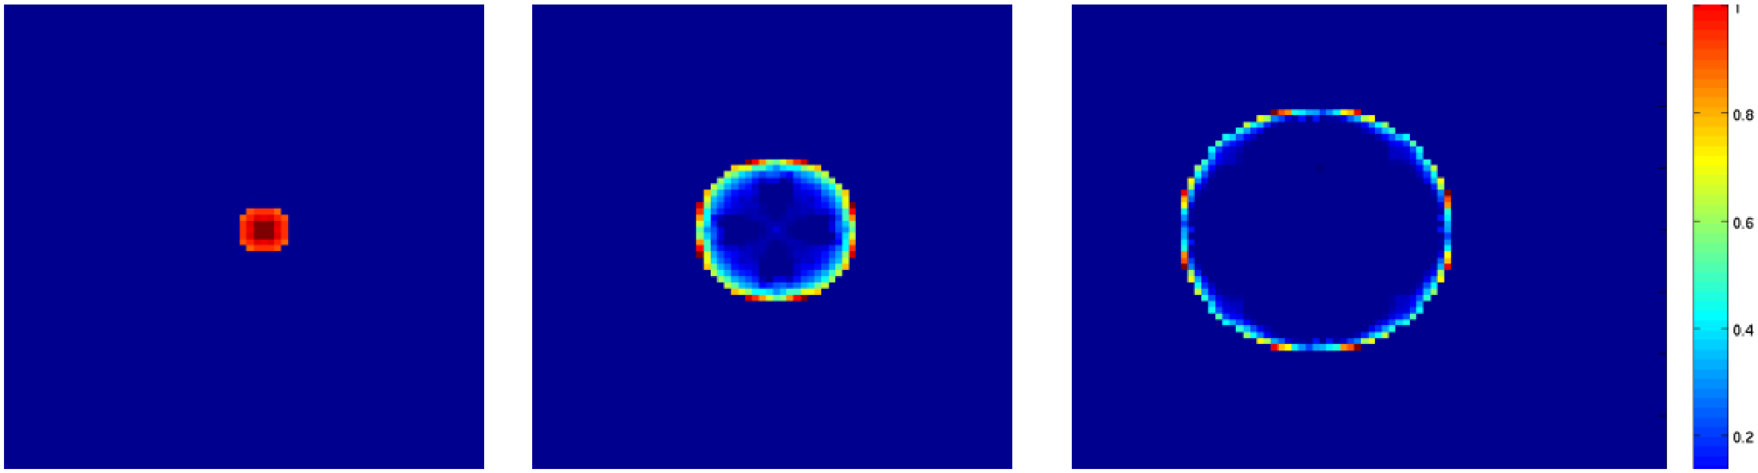

Model is OK on given segments. But, how do we optimize non-submodular 3-cliques (010) and (101)?

1. Standard trick: convert to non-submodular pair-wise binary optimization

2. Our observation: ~~QPBO~~ does not work (unless non-submodular regularization is very weak)

Fast Trust Region [CVPR13, arXiv]

uses *local submodular approximations*

Segmentation Examples

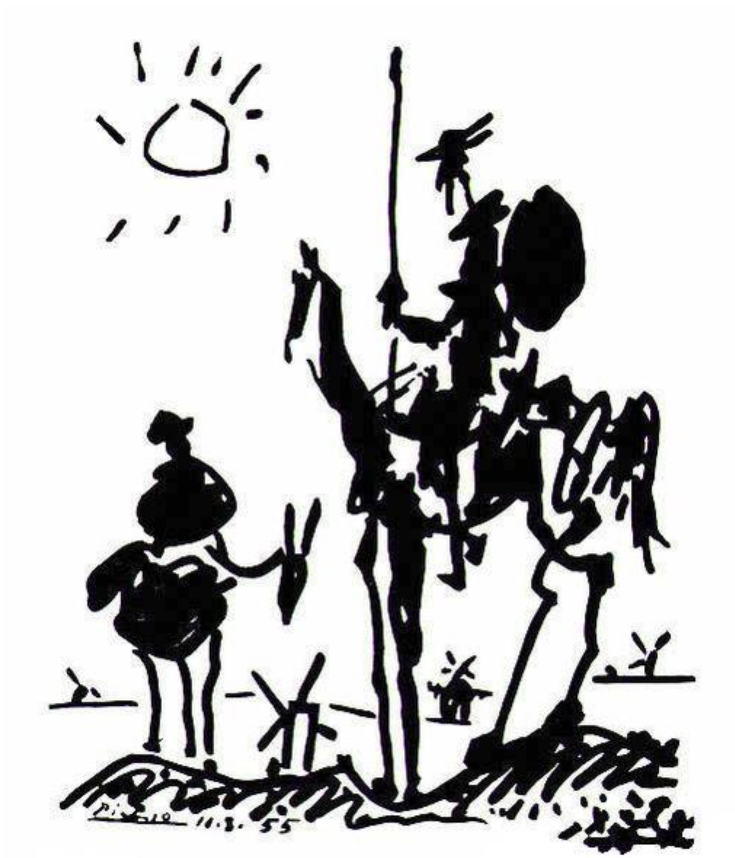

length-based regularization

Segmentation Examples

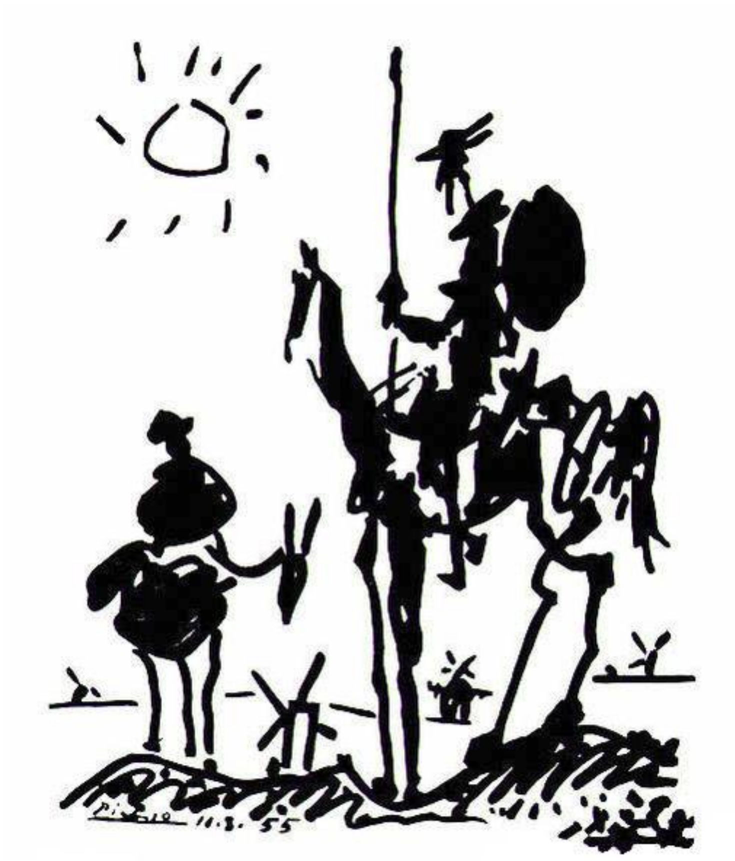

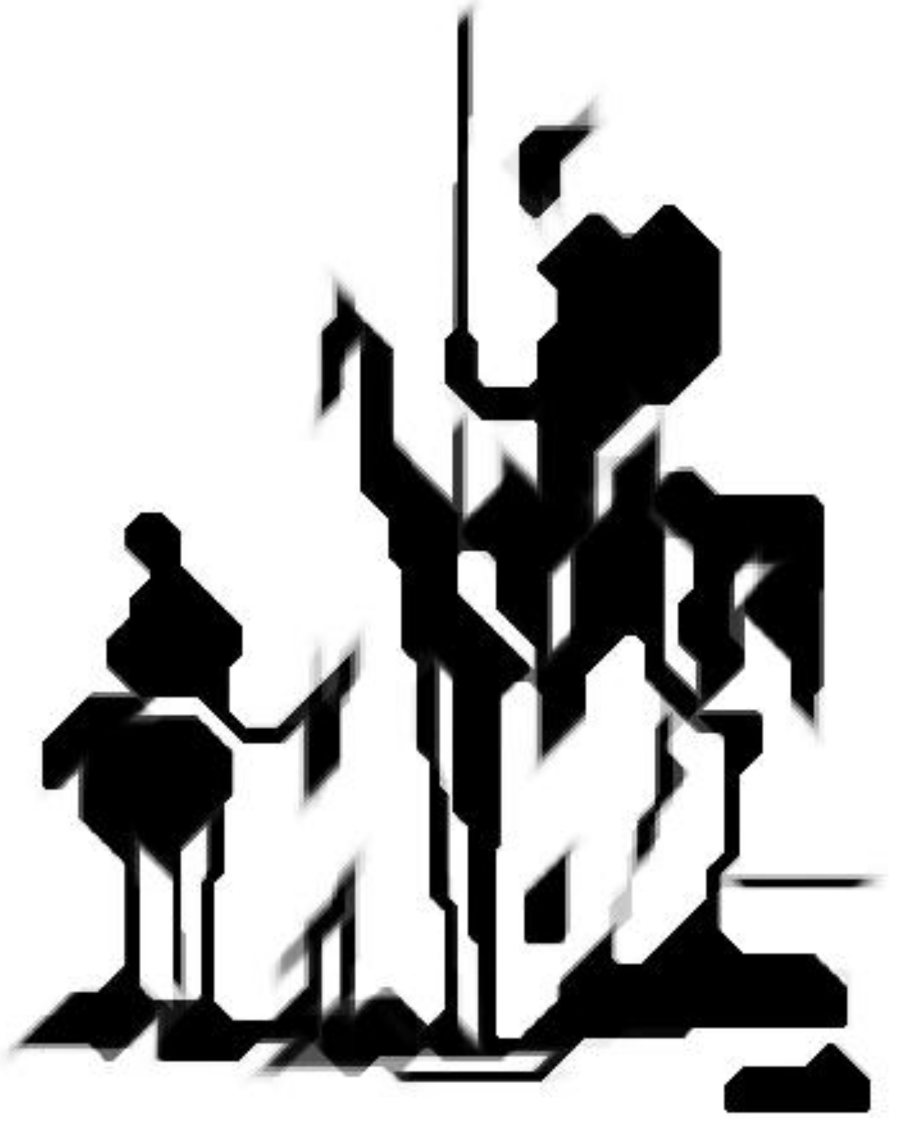

elastica [Heber,Ranftl,Pock, 2012]

Segmentation Examples

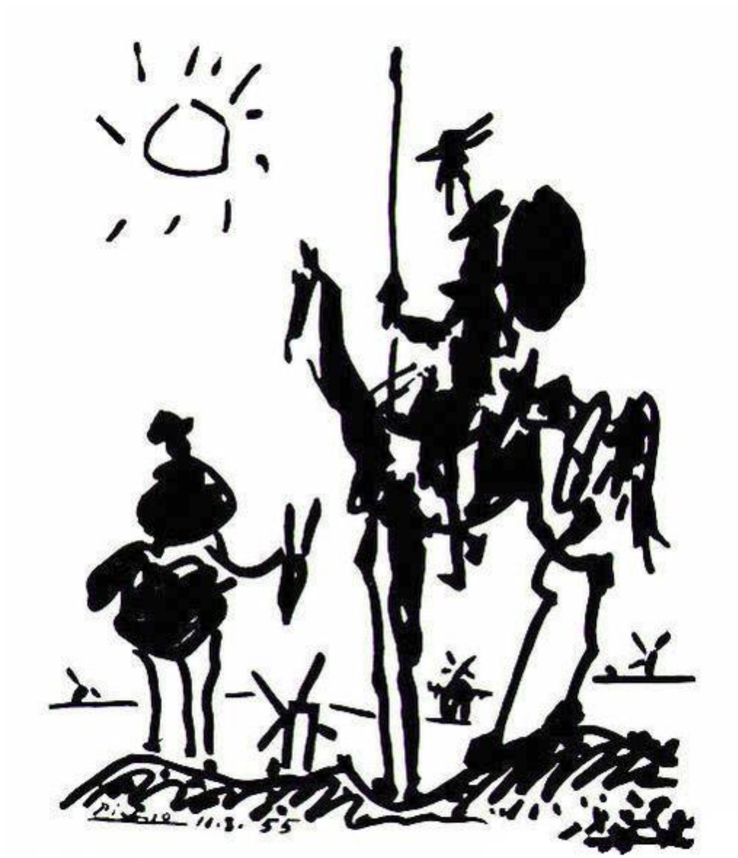

90-degree curvature [El-Zehiry&Grady, 2010]

Segmentation Examples

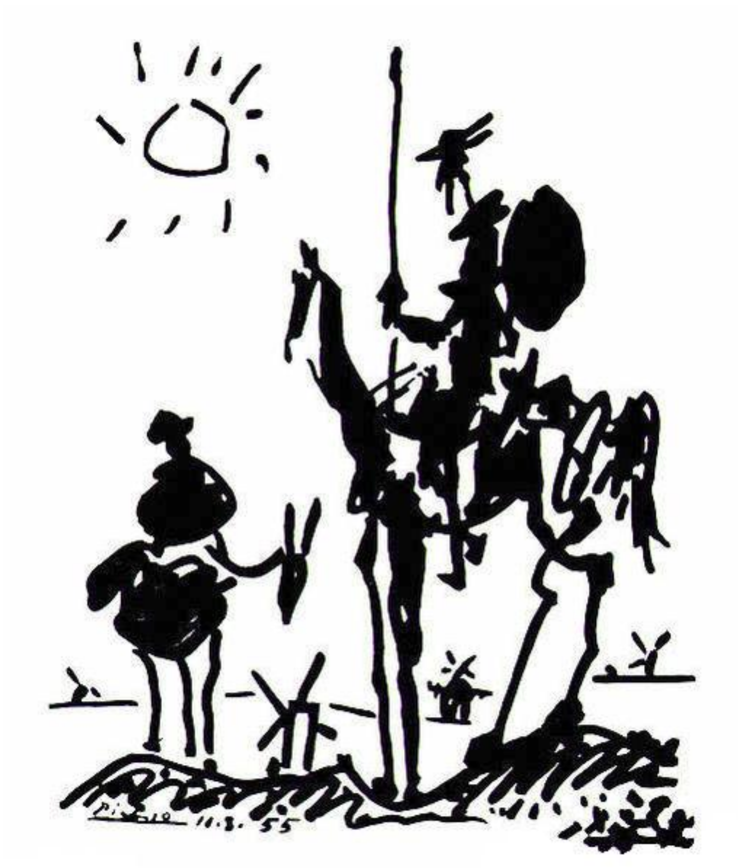

7x7 neighborhood

our squared curvature

Segmentation Examples

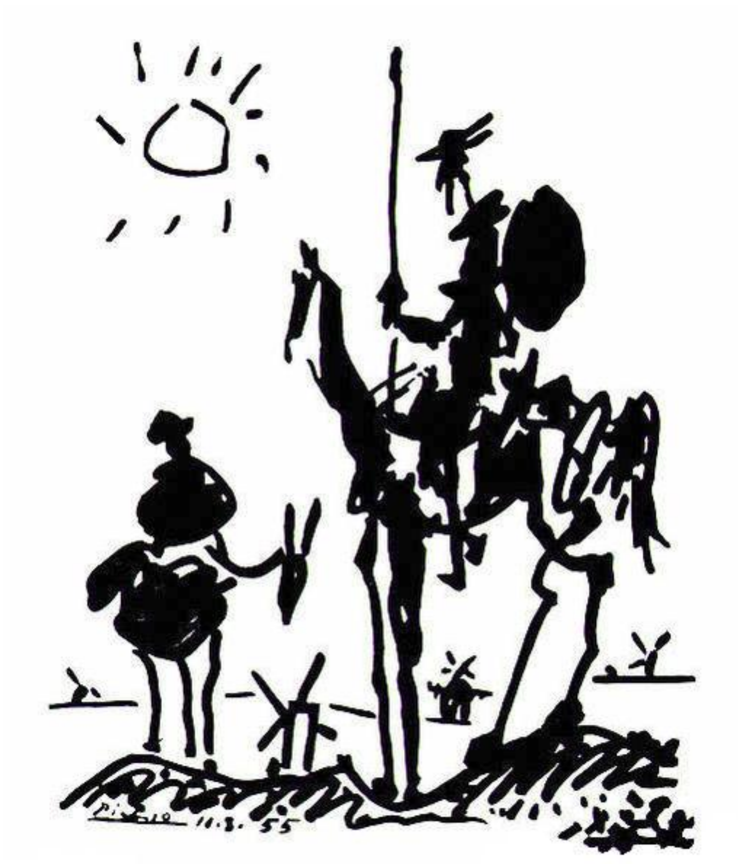

7x7 neighborhood

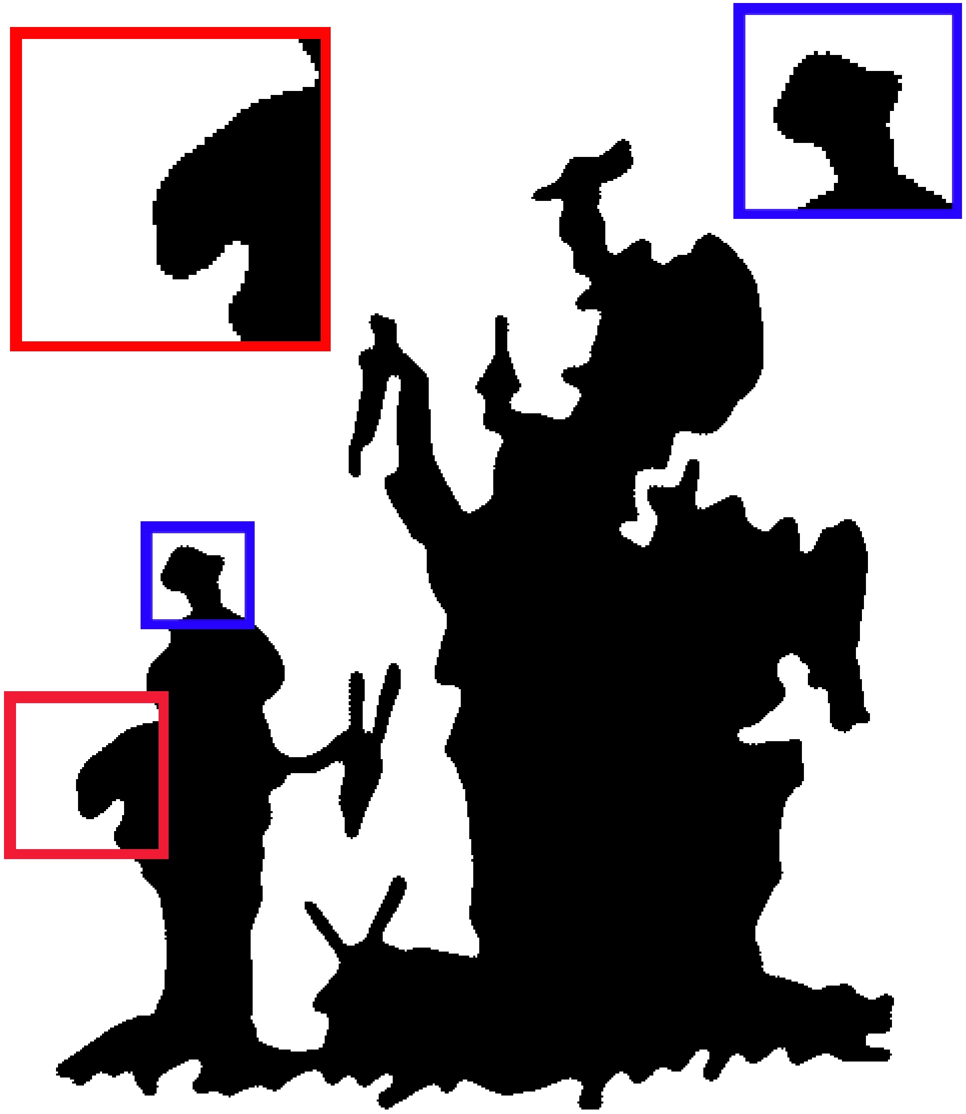

our squared curvature (stronger)

Segmentation Examples

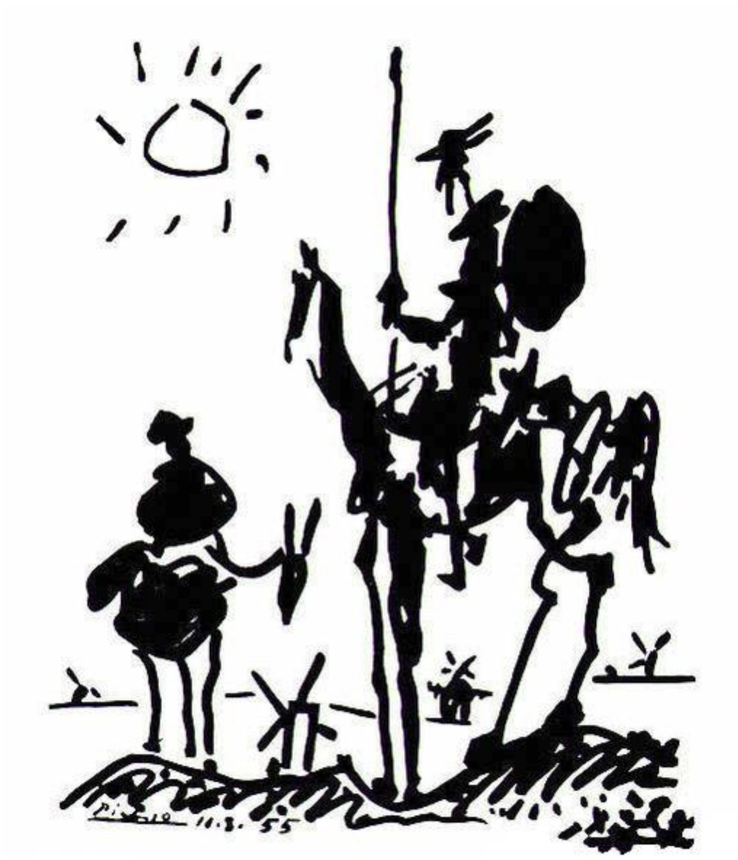

2x2 neighborhood

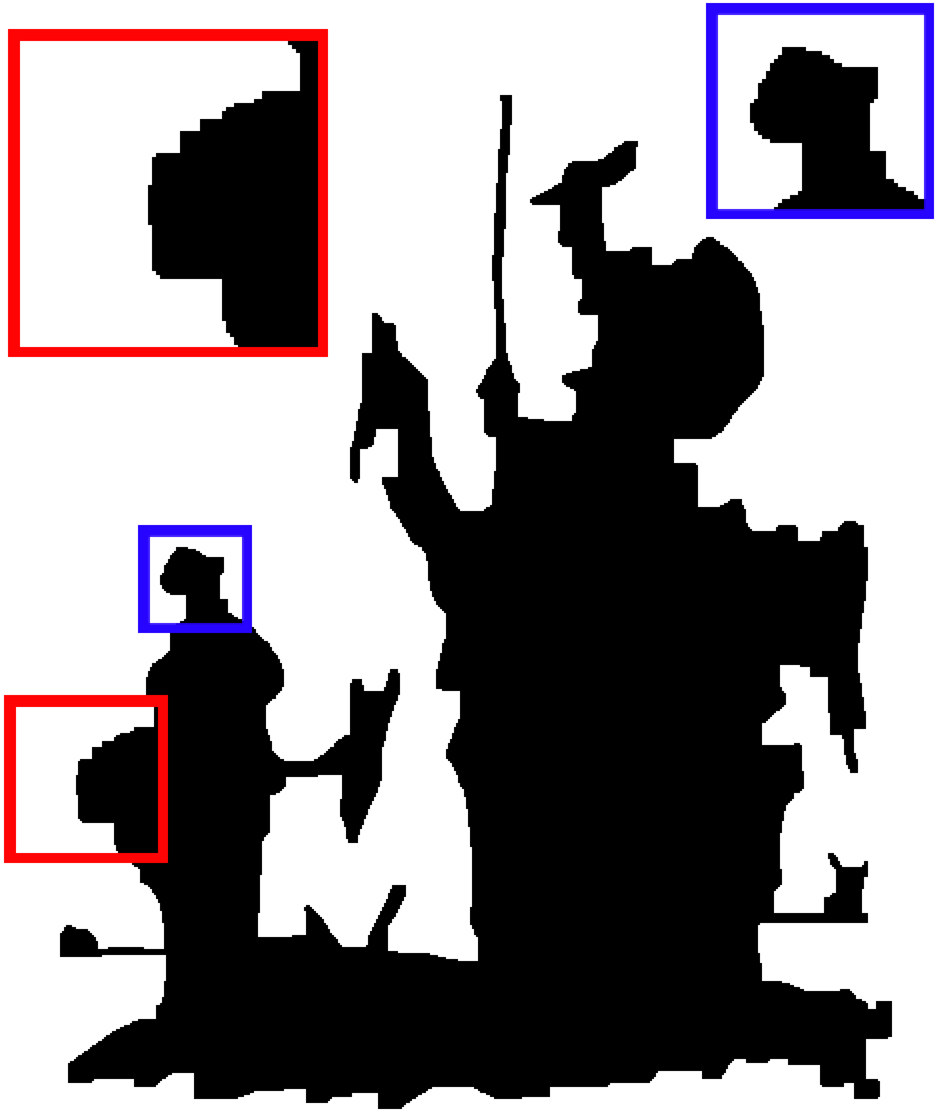

our squared curvature (stronger)

Binary inpainting

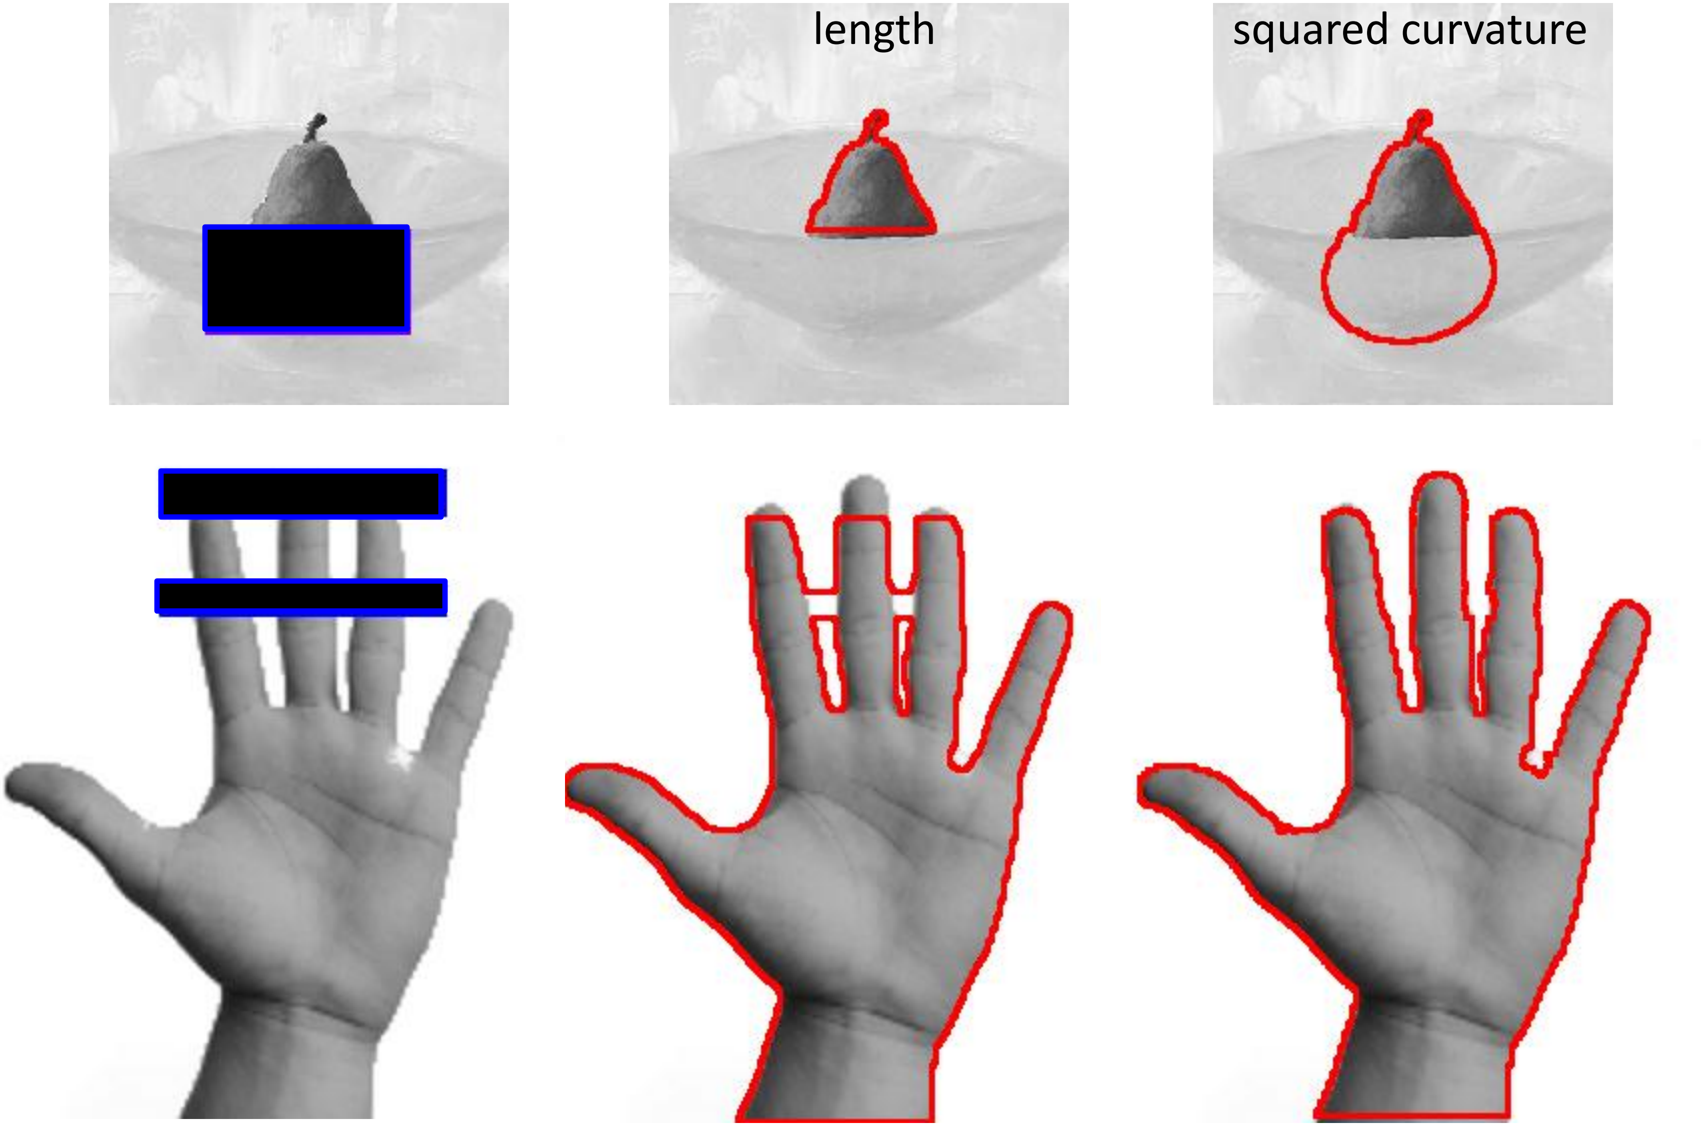

Conclusions

- **Optimization of Entropy** is a useful information-theoretic interpretation of color model estimation
- **L_1 color separation** is an easy-to-optimize objective useful in its own right [ICCV 2013]
- **Global optimization matters**: one cut [ICCV13]
- **Trust region, auxiliary cuts, partial enumeration**

General approximation techniques

- for high-order energies [CVPR13]

- for non-submodular energies [arXiv'13]

outperforming state-of-the-art combinatorial optimization methods